GSDC
Global Skill Development Council

GENERATIVE AI FOR LEADERS 50+ PRACTICE MCQS

www.gsdcouncil.org

DOMAIN 1: Foundations of Generative AI — 20 Questions

Q1

What is the most accurate definition of Generative AI?

- A) AI that classifies existing data into predefined categories
- B) AI that creates new content by learning patterns from training data
- C) AI that monitors and detects anomalies in datasets
- D) AI that automates rule-based repetitive tasks

✅ **Answer: B** — Generative AI produces new content — text, images, code, audio, video — by learning statistical patterns from large training datasets. It differs from discriminative AI (which classifies) and robotic process automation (which follows rules). The "generative" distinction is its ability to create, not just categorize.

Q2

Which of the following best describes a Large Language Model (LLM)?

- A) A database that stores large volumes of structured text
- B) A rule-based system that matches queries to pre-written responses
- C) A neural network trained on massive text data to predict and generate language
- D) A search engine that indexes and retrieves web content

✅ **Answer: C** — LLMs are deep learning models — specifically transformer-based neural networks — trained on billions of words of text. They learn to predict the next token in a sequence, which enables them to generate coherent, contextually relevant language. They do not search the internet or follow pre-written rules.

Q3

What is a "token" in the context of Large Language Models?

- A) A security credential used to authenticate API access
- B) A unit of text — roughly a word or part of a word — that LLMs process
- C) A monetary unit used to purchase AI compute credits
- D) A tagged label applied to training data during supervised learning

✅ **Answer: B** — Tokens are the basic units LLMs work with. A token is approximately ¾ of a word in English — so 100 tokens ≈ 75 words. Token limits define how much text a model can process in one session (its context window). Understanding tokens matters for leaders because API pricing is token-based and context window size affects use case design.

Q4

What does the term "context window" refer to in Gen AI systems?

- A) The graphical user interface through which users interact with an AI model
- B) The maximum amount of text an LLM can process and remember in a single interaction
- C) The time period during which a model's training data was collected
- D) The browser window used to access AI tools

✅ **Answer: B** — The context window is the maximum number of tokens an LLM can "see" at once — input plus output combined. Larger context windows (e.g., Gemini 1.5 Pro's 1 million tokens) allow processing of entire codebases or lengthy documents. This directly impacts which enterprise use cases are feasible — a small context window limits document analysis tasks.

Q5

What is the primary difference between a Foundation Model and a fine-tuned model?

- A) Foundation models are open-source; fine-tuned models are proprietary
- B) Foundation models are general-purpose pre-trained models; fine-tuned models are adapted to specific tasks or domains
- C) Foundation models require more compute at inference; fine-tuned models are faster
- D) Foundation models only process text; fine-tuned models handle images

✅ **Answer: B** — Foundation models (e.g., GPT-4, Claude, Gemini) are trained on broad data and capable across many tasks. Fine-tuning adapts a foundation model to a specific domain, task, or organizational context using additional targeted training data. Fine-tuning improves performance on specific tasks but requires data, compute, and ML expertise — important considerations for leaders evaluating build vs. buy decisions.

Q6

A leader asks: "Our AI tool keeps confidently stating wrong facts. What is this problem called?"

- A) Overfitting
- B) Model drift
- C) Hallucination
- D) Underfitting

✅ **Answer: C** — Hallucination occurs when an LLM generates plausible-sounding but factually incorrect information — and presents it with confidence. It is a fundamental limitation of LLMs, not a bug. Mitigation strategies include RAG (grounding responses in verified data), human-in-the-loop review, and confidence scoring. Leaders must build workflows that account for hallucination risk, especially in regulated domains.

Q7

What does RAG stand for, and what problem does it solve?

- A) Recursive Attention Generation — it improves model speed
- B) Retrieval-Augmented Generation — it grounds AI responses in real-time, verified information
- C) Random Approximate Grouping — it clusters training data more efficiently
- D) Regulated AI Governance — it ensures compliance in enterprise deployments

✅ **Answer: B** — RAG connects an LLM to an external knowledge base — your company's documents, databases, or the internet. Before generating a response, the system retrieves relevant documents and passes them to the model as context. This dramatically reduces hallucinations, keeps responses current, and allows AI to work with proprietary organizational knowledge without retraining the model.

Q8

Which parameter controls how creative or random an LLM's outputs are?

- A) Learning rate
- B) Batch size
- C) Temperature
- D) Dropout rate

✅ **Answer: C** — Temperature controls the probability distribution of the model's next-token predictions. Low temperature (near 0) = deterministic, predictable, consistent outputs — ideal for factual tasks like data extraction. High temperature (near 1 or above) = more creative, varied, surprising outputs — better for brainstorming or creative writing. Leaders configuring AI tools for different functions should understand this parameter.

Q9

What is the key difference between zero-shot and few-shot prompting?

- A) Zero-shot uses no AI model; few-shot uses a small model
- B) Zero-shot provides no examples in the prompt; few-shot includes 2–5 examples to guide the model
- C) Zero-shot is faster computationally; few-shot is more accurate but slower
- D) Zero-shot works only with text; few-shot works with images

✅ **Answer: B** — Zero-shot prompting asks the model to perform a task without any examples — relying entirely on its training. Few-shot prompting includes examples of the desired input-output pattern directly in the prompt, significantly improving output quality and format consistency. Leaders designing AI workflows should default to few-shot approaches for tasks requiring specific output formats or domain-specific language.

Q10

What is "prompt engineering" and why does it matter for leaders?

- A) The technical process of training an AI model from scratch
- B) The practice of designing effective inputs to guide AI output quality, format, and relevance
- C) The engineering discipline responsible for building AI infrastructure
- D) A method for encrypting prompts to protect intellectual property

✅ **Answer: B** — Prompt engineering is the art and science of crafting inputs that reliably produce high-quality AI outputs. For leaders, it is a strategic skill — the difference between AI that produces generic outputs and AI that delivers precise, business-ready results. Effective prompts include context, constraints, format instructions, and examples. This skill does not require coding and is immediately applicable at every leadership level.

Q11

Which of the following best describes a "multimodal AI" model?

- A) A model that can be deployed across multiple cloud platforms simultaneously
- B) A model that processes and generates across multiple content types — text, images, audio, and video
- C) A model trained on data from multiple industries simultaneously
- D) A model that supports multiple languages in the same session

✅ **Answer: B** — Multimodal AI integrates multiple data modalities in a single model. GPT-4o, Gemini, and Claude 3 can all process both text and images. This unlocks use cases like analyzing a chart image and explaining it in text, or processing a document scan and extracting structured data. Leaders should identify workflows involving multiple content types as prime multimodal AI opportunities.

Q12

What is the transformer architecture's primary innovation that enabled the rise of modern LLMs?

- A) Convolutional layers that process text as visual data
- B) Recurrent cells that process text sequentially with memory
- C) An attention mechanism that weighs the relevance of every word in a sequence to every other word simultaneously
- D) A gradient boosting approach that combines multiple smaller models

✅ **Answer: C** — The transformer's "self-attention mechanism" (introduced in the 2017 paper "Attention Is All You Need") allows models to process all words in a sequence simultaneously — understanding relationships between distant words far more effectively than previous architectures. This breakthrough made the development of GPT, BERT, and all modern LLMs possible. Leaders don't need to understand the math, but knowing why transformers matter explains why AI capabilities improved so dramatically after 2017.

What distinguishes an "AI Agent" from a standard LLM chatbot?

- A) AI agents are more accurate because they use larger datasets
- B) AI agents can plan multi-step tasks, use external tools, and take actions autonomously — not just generate text
- C) AI agents require human approval for every output they produce
- D) AI agents only work with structured data, not natural language

✅ **Answer: B** — A chatbot generates responses. An AI agent plans, reasons, executes, and iterates. Agents can browse the web, write and run code, call APIs, manage files, and chain together multi-step workflows — often with minimal human intervention. For leaders, agentic AI represents the next wave of enterprise automation — moving from AI that assists humans to AI that acts on behalf of humans within defined boundaries.

What is "model fine-tuning" and when is it appropriate for enterprises?

- A) Adjusting the model's user interface for better usability
- B) Further training a pre-trained foundation model on domain-specific data to improve performance on targeted tasks
- C) Reducing a model's size to improve inference speed in production
- D) Updating a model's training data with real-time internet content

✅ **Answer: B** — Fine-tuning adds a training step where the model learns from your specific data — improving performance on domain-specific tasks, adopting your brand's voice, or learning industry-specific terminology. It requires labeled training data, compute resources, and ML expertise. For most enterprise use cases, RAG and prompt engineering deliver sufficient results at far lower cost and complexity. Fine-tuning is appropriate when consistent, specialized behavior is critical and prompt-based approaches are insufficient.

Q15

What are "embeddings" in the context of Gen AI?

- A) Watermarks embedded in AI-generated content for copyright protection
- B) Numerical vector representations of text that capture semantic meaning for similarity search
- C) Hyperlinks embedded in AI-generated documents
- D) Hidden instructions embedded in prompts to guide model behavior

✅ **Answer: B** — Embeddings convert text into numerical vectors in a high-dimensional space, where semantically similar content is positioned close together. They power semantic search (finding conceptually related content even when exact words differ), recommendation systems, and RAG retrieval. For leaders, embeddings are the engine behind "ask questions of your documents" use cases — they allow AI to find relevant content from large knowledge bases quickly and accurately.

Q16

What is the difference between "open-source" and "proprietary" AI models — and why does it matter strategically?

- A) Open-source models are less accurate; proprietary models are more capable in all situations
- B) Open-source models can be self-hosted and modified; proprietary models are accessed via API with vendor control
- C) Open-source models are free to use commercially without any restrictions
- D) Proprietary models are always safer because they have more governance oversight

✅ **Answer: B** — Open-source models (Meta's Llama, Mistral) can be downloaded, self-hosted, and customized — offering data privacy, no ongoing API costs, and freedom from vendor lock-in. Proprietary models (GPT-4, Claude, Gemini) are accessed via API — faster to deploy, typically more capable out-of-the-box, but create vendor dependency and data-sharing considerations. Strategic leaders evaluate this trade-off based on data sensitivity, technical capability, and long-term cost modeling.

Q17

What is "model drift" and why should leaders monitor for it?

- A) The gradual migration of AI workloads from one cloud provider to another
- B) The degradation of a deployed AI model's performance over time as real-world data diverges from training data
- C) The tendency of AI models to produce increasingly biased outputs the longer they run
- D) A security vulnerability where models are manipulated through malicious prompts

✅ **Answer: B** — Model drift occurs when the distribution of real-world data changes after deployment, causing the model's outputs to become less accurate or relevant. A fraud detection model trained on pre-pandemic transactions may perform poorly as spending patterns change. Leaders must ensure AI governance includes ongoing model performance monitoring, not just deployment — AI is not a "set and forget" technology.

Q18

What is "prompt injection" and why is it a security concern for leaders?

- A) A technique for injecting training data into a model post-deployment
- B) A malicious technique where hidden instructions in content override an AI system's intended behavior
- C) A method for compressing long prompts to reduce API token costs
- D) An approach for combining multiple prompts to improve output quality

✅ **Answer: B** — Prompt injection occurs when malicious instructions hidden in content the AI processes (e.g., a document, email, or webpage) override the system's intended instructions. For example, a webpage could contain hidden text telling an AI agent to leak information or take unauthorized actions. As organizations deploy AI agents that process external content, prompt injection becomes a serious security risk requiring technical safeguards and policy controls.

Q19

What does "RLHF" stand for and why is it significant?

- A) Reinforcement Learning from Human Feedback — the training technique used to align LLMs with human preferences and values
- B) Real-time Language Handling Framework — a method for processing streaming AI outputs
- C) Regulated Language Harmonization Function — a compliance standard for enterprise AI
- D) Recursive Loop Handling Framework — a technique for managing multi-turn conversations

✅ **Answer: A** — RLHF is how models like ChatGPT and Claude are made helpful, harmless, and honest. Human raters evaluate model outputs, and the model is trained to produce responses that humans prefer. This is the key technique that transformed raw LLMs into useful assistants — and it is why modern AI models are generally more cautious, more helpful, and less likely to produce harmful content than their predecessor models. Leaders should understand this when evaluating AI safety claims from vendors.

A Gen AI system produces outputs that systematically disadvantage candidates from certain demographic groups in hiring decisions. What is this problem called?

- A) Model hallucination
- B) Data poisoning
- C) Algorithmic bias
- D) Context window overflow

✅ **Answer: C** — Algorithmic bias occurs when an AI system produces outputs that systematically

favor or disadvantage certain groups — typically because the training data reflected historical human biases. In hiring, an AI trained on past hiring decisions may perpetuate patterns of discrimination. Leaders must implement regular bias audits, diverse training data practices, and human oversight at high-stakes decision points to mitigate this risk.

Key Domain 1 Concepts

- Tokens & Context Windows
- Foundation Models vs. Fine-Tuning
- RAG & Embeddings

Critical Risks to Know

- Hallucination
- Model Drift
- Prompt Injection
- Algorithmic Bias

Key Techniques

- Zero-shot vs. Few-shot Prompting
- Temperature Control
- RLHF Alignment

Q21

Which framework is most useful for a leader evaluating whether to build, buy, or partner for AI capabilities?

- A) Evaluate solely on upfront cost — choose whichever option is cheapest initially
- B) Assess competitive differentiation, data sensitivity, time-to-value, and in-house talent availability
- C) Always buy from established vendors to minimize risk
- D) Build all AI capabilities internally to maintain control

✅ **Answer: B** — The build-buy-partner decision requires multi-dimensional analysis. Build when AI is core to competitive differentiation and talent is available. Buy when speed matters and the use case is non-differentiating. Partner when domain expertise must combine with technology. No single option is universally correct — leaders must evaluate each initiative individually against these criteria.

Q22

What is an AI Center of Excellence (CoE) and what is its primary function?

- A) A physical AI research lab where models are developed from scratch
- B) A centralized team that sets standards, governance, and best practices while enabling AI adoption across business units
- C) An external consulting team hired to manage all AI projects
- D) A regulatory body that approves AI deployments before they go live

✅ **Answer: B** — An AI CoE provides centralized coordination without centralizing all AI work. It creates reusable frameworks, governance policies, shared platforms, and capability-building programs. Business units still own their AI use cases, but the CoE prevents duplication, ensures consistency, manages risk, and accelerates learning across the organization. It is typically the most effective structural model for enterprise AI scaling.

Q23

A leader wants to prioritize which AI use cases to pursue first. Which framework is most appropriate?

- A) Alphabetical order by department name
- B) A 2×2 matrix plotting Business Value against Implementation Complexity
- C) Select the most technically impressive use cases to demonstrate AI capability
- D) Start with the most complex use cases to build organizational capability fastest

✅ **Answer: B** — The Value vs. Complexity matrix is the standard prioritization framework. High Value + Low Complexity = Quick wins that build momentum and demonstrate ROI. High Value + High Complexity = Strategic bets requiring careful roadmap planning. Starting with quick wins is critical — early proof of value secures continued investment and builds organizational confidence before tackling complex initiatives.

What is "shadow AI" and what is the appropriate leadership response?

- A) AI tools running on backup servers when primary systems are down — respond by improving infrastructure resilience
- B) Employees using unauthorized AI tools for work tasks — respond with enablement, approved alternatives, and clear policies rather than blanket prohibition
- C) AI systems that operate without any user interface — respond by adding human oversight layers
- D) Competitor AI systems that replicate your products — respond with intellectual property protections

> ✅ **Answer: B** — Shadow AI refers to the widespread use of consumer AI tools (personal ChatGPT accounts, etc.) for work purposes — often involving the upload of confidential data to unsanctioned platforms. Surveys indicate 40–60% of employees are already doing this. Prohibition doesn't work — it just drives the behavior underground. Effective responses include providing approved AI tools that meet employee needs, establishing clear acceptable use policies, implementing data loss prevention controls, and communicating the risks transparently.

Which of the following best describes a "responsible AI" framework?

- A) A framework that prohibits AI use in any high-risk business decision
- B) A set of principles and practices ensuring AI is developed and deployed in a fair, accountable, transparent, and safe manner
- C) A vendor certification program that guarantees AI tools are free from bias
- D) A legal compliance program focused exclusively on GDPR adherence

> ✅ **Answer: B** — Responsible AI frameworks typically encompass fairness (bias detection and mitigation), accountability (clear ownership of AI decisions), transparency (explainable outputs), safety (fail-safes and human oversight), and privacy (data governance). They are operationalized through governance boards, bias audits, model documentation, and human-in-the-loop requirements — not through avoidance of AI in complex domains.

Q26

What does "Human-in-the-Loop" (HITL) mean in an enterprise AI context?

- A) Requiring a human to type inputs into AI systems rather than using automated feeds
- B) Maintaining meaningful human oversight and decision authority at critical points in AI-assisted workflows
- C) A user interface design principle ensuring AI tools are intuitive for human users
- D) A legal requirement that all AI outputs must be reviewed by lawyers before use

✅ **Answer: B** — HITL is the practice of keeping human judgment in the decision chain at high-stakes moments — even when AI handles the analysis. For example, an AI might screen resumes and rank candidates, but a human makes the final interview decision. HITL requirements should scale with decision risk: low-risk decisions may need minimal oversight, while healthcare, legal, or financial decisions typically require mandatory human review.

A company discovers that an AI system it deployed 18 months ago is now performing significantly worse than at launch. What is the most likely explanation?

- A) The AI model has become "bored" and is generating random outputs
- B) Model drift — the real-world data distribution has changed since the model was trained, degrading its accuracy
- C) The model has been hacked and is deliberately underperforming
- D) All AI models are programmed to degrade after 12–18 months to encourage upgrades

✅ **Answer: B** — Model drift is a natural occurrence as the world changes and real-world data diverges from training data. A demand forecasting model trained before a major market shift will underperform post-shift. Leaders must budget for ongoing model monitoring, retraining schedules, and performance benchmarking as part of any AI deployment — not just initial deployment costs.

What is the most critical success factor in enterprise AI transformation, according to research?

- A) The sophistication of the AI models deployed
- B) The size of the AI budget allocated
- C) Change management — addressing people, culture, and adoption challenges
- D) The cloud infrastructure supporting AI workloads

✅ **Answer: C** — The majority of AI transformation failures are attributable to people and culture challenges — not technology limitations. Resistance to adoption, fear of job displacement, lack of AI literacy, and poor communication of the AI vision consistently outrank technical factors as reasons AI initiatives stall. Leaders who invest proportionally in change management alongside technology deployment dramatically improve success rates.

Which organizational structure is most effective for scaling AI across a large enterprise?

- A) Fully centralized — all AI work done by a single central team
- B) Fully decentralized — each business unit independently develops its own AI capabilities
- C) Federated hub-and-spoke — central CoE sets standards and platforms while business units own domain-specific implementations
- D) Outsourced — all AI development and deployment handled by a single strategic vendor

✅ **Answer: C** — The hub-and-spoke (federated) model balances governance with agility. A central CoE prevents duplication, enforces ethical standards, manages shared infrastructure, and transfers knowledge. Business unit-embedded AI leads ensure solutions are domain-relevant and adopted. Fully centralized models create bottlenecks; fully decentralized models produce inconsistency and governance gaps; full outsourcing creates dangerous capability dependency.

What should be included in an organization's AI Acceptable Use Policy?

- A) Only technical specifications for approved AI tools
- B) Approved tools, prohibited use cases, data handling rules, disclosure requirements, and consequences for policy violations
- C) A complete prohibition on AI use in any customer-facing process
- D) Only compliance requirements for GDPR and CCPA

✅ **Answer: B** — A comprehensive AI Acceptable Use Policy covers: which tools are approved and how to access them, what data can and cannot be used with AI (data classification guidance), use cases that are prohibited (e.g., fully autonomous HR decisions), requirements to disclose AI-generated content in certain contexts, and the consequences of policy violations. The policy must be practical — overly restrictive policies accelerate shadow AI adoption.

What is the primary risk of deploying AI in high-stakes decisions without explainability?

- A) The AI will take longer to process decisions
- B) Leaders cannot justify, audit, or defend decisions — creating legal, ethical, and reputational risk
- C) Employees will distrust AI more than necessary
- D) The AI will produce less accurate outputs without explanation requirements

✅ **Answer: B** — Explainability (XAI — Explainable AI) is critical when AI influences consequential decisions affecting people's lives — credit, hiring, healthcare, legal. If an AI denies someone a loan and neither the institution nor the individual can understand why, this creates regulatory exposure, discrimination risk, and destroys trust. Leaders in regulated industries must demand explainability from AI vendors as a non-negotiable requirement.

Q32

A leader is asked to evaluate an AI vendor's data privacy practices. Which question is MOST important to ask?

- A) "How many customers do you currently have?"
- B) "Will our data be used to train your models, and can we opt out of that?"
- C) "What is your net promoter score?"
- D) "How many GPUs do you have in your data center?"

✅ **Answer: B** — One of the most critical data privacy questions for enterprise AI is whether your inputs are used to train the vendor's future models — a common practice with consumer AI tools that creates significant confidentiality risks. Enterprise contracts should explicitly prohibit training on customer data and should be backed by technical controls, not just policy. This question is foundational to any AI vendor due diligence process.

Q33

What does "AI literacy" mean for a non-technical business leader?

- A) The ability to write code in Python or another programming language
- B) Understanding AI capabilities, limitations, vocabulary, and strategic implications well enough to make informed decisions
- C) Completing an advanced machine learning certification course
- D) Reading all academic papers published on AI each month

✅ **Answer: B** — AI literacy for leaders is not about technical implementation — it is about informed decision-making. A literate AI leader understands what AI can and cannot do, knows the vocabulary to engage with technical teams and vendors meaningfully, can spot AI hype vs. genuine capability, and has a framework for evaluating AI investments and risks. This is why certification programs like GSDC's Generative AI for Leaders are designed for business leaders, not engineers.

Q34

What is the "pilot-to-scale" framework in enterprise AI deployment?

- A) Deploying AI at full scale immediately to maximize first-mover advantage
- B) Starting with a narrow, validated proof of concept before investing in full-scale deployment — reducing risk and building evidence for broader investment
- C) Piloting AI with external customers before deploying internally
- D) Testing AI with technical teams only and never scaling to business users

✅ **Answer: B** — The pilot-to-scale approach manages AI investment risk. A well-designed pilot: has a clearly defined use case, measurable success criteria, a realistic timeline (typically 30–90 days), and a pre-agreed decision framework for scaling or stopping. This approach builds organizational confidence, generates internal case studies, and surfaces technical and adoption challenges at manageable scale before they become enterprise-wide problems.

Q35

Which of the following is NOT typically a metric for measuring AI program success?

- A) Reduction in task completion time
- B) Employee AI tool adoption rate
- C) Number of AI models deployed
- D) Customer satisfaction score improvement

✅ **Answer: C** — The number of models deployed is a vanity metric — it measures activity, not value. Meaningful AI program metrics fall into three categories: operational (time saved, error reduction, cost per task), business (revenue impact, customer satisfaction, speed to market), and adoption (active users, task completion rates, engagement). Leaders who optimize for model count rather than business outcomes consistently fail to demonstrate AI ROI.

What is "AI governance" and what does it encompass in an enterprise context?

- A) The IT department's management of AI software licenses and infrastructure
- B) The policies, processes, roles, and oversight mechanisms that ensure AI is developed and used responsibly, effectively, and in alignment with organizational values
- C) A government regulatory framework that organizations must comply with
- D) The process by which AI models are audited for technical accuracy

✅ **Answer: B** — AI governance is broader than compliance. It includes: defining AI principles and values, establishing review and approval processes for AI deployments, creating accountability structures (who owns AI decisions), managing AI risk registries, ensuring ongoing monitoring of deployed systems, and creating feedback mechanisms for identifying and correcting problems. It is a strategic leadership function, not just an IT or legal concern.

Q37

A newly appointed Chief AI Officer asks for the single highest-priority action in their first 90 days. What should it be?

- A) Deploy as many AI tools as possible to demonstrate momentum
- B) Conduct a comprehensive AI readiness assessment across people, process, data, and infrastructure
- C) Replace all existing technology with AI-native alternatives
- D) Hire an external AI consulting firm to manage all AI strategy

✅ **Answer: B** — Before setting strategy or deploying tools, a leader needs an honest picture of organizational reality. An AI readiness assessment across the four pillars — people (AI literacy, talent), process (workflows suitable for AI augmentation), data (quality, governance, accessibility), and infrastructure (cloud maturity, API capabilities) — provides the diagnostic foundation for all subsequent decisions. Acting without this assessment leads to misaligned investments and failed pilots.

Q38

What is the most significant risk of creating a fully centralized AI team with no embedded AI capability in business units?

- A) The central team will use too much computing budget
- B) AI solutions will be technically excellent but poorly adopted because they are disconnected from business context and user needs
- C) The central team will develop AI that is too powerful for business users
- D) Regulatory authorities will prohibit centralized AI governance

✅ **Answer: B** — The "ivory tower" failure mode in AI transformation: a highly capable central team builds sophisticated AI tools that business units don't use because they don't understand them, don't trust them, or they don't fit into real workflows. Embedding AI capability in business units — even just one "AI Champion" per department — creates the contextual understanding and change management bridge that translates AI capability into actual adoption.

What does Total Cost of Ownership (TCO) include for an enterprise AI program beyond licensing fees?

- A) Only the monthly SaaS subscription cost
- B) Compute costs, data preparation, integration engineering, training programs, human oversight roles, ongoing monitoring, and compliance review
- C) Only the cost of AI hardware and infrastructure
- D) Only the salaries of the AI development team

✅ **Answer: B** — Leaders consistently underestimate AI TCO. The subscription or API cost is often the smallest component. Data preparation and cleaning typically consumes 60–70% of project effort. Integration with existing systems requires engineering resources. Change management and training programs are substantial. Ongoing human oversight, quality assurance, model monitoring, and compliance review are recurring costs. A complete TCO model is essential before any AI investment decision.

Which of the following best describes an "AI use case" that should NOT be pursued, regardless of technical feasibility?

- A) A use case where humans can be removed from the process entirely
- B) A use case that automates a high-volume, low-value repetitive task
- C) A use case that makes fully autonomous, unexplained decisions affecting an individual's employment, credit, or healthcare without human oversight
- D) A use case that costs more to implement than it saves in year one

✅ **Answer: C** — Technical feasibility does not determine ethical permissibility. Fully autonomous AI decisions affecting fundamental aspects of people's lives — with no explainability and no human oversight — violate responsible AI principles and, in many jurisdictions, regulatory requirements (EU AI Act classifies these as high-risk AI systems). Leaders must distinguish between what AI can do and what it should do in each organizational context.

What is the best approach to managing employee fear and resistance during an AI transformation?

- A) Minimize communication about AI initiatives to prevent anxiety
- B) Acknowledge concerns directly, communicate transparently about intentions, invest in reskilling, involve employees in design, and show early wins that benefit the team
- C) Replace resistant employees first to accelerate cultural change
- D) Frame all AI as "productivity tools" and avoid mentioning automation

✅ **Answer: B** — Research on successful change management consistently shows that transparency and involvement are more effective than minimization or avoidance. Employees who feel informed, consulted, and equipped for change are significantly more likely to adopt new technologies. Leaders who acknowledge legitimate concerns, provide concrete reskilling commitments, and demonstrate early AI benefits for employees — not just for shareholders — build the trust necessary for sustainable adoption.

A retail company wants to use AI to optimize its inventory management. What is the FIRST step a leader should take?

- A) Purchase the most advanced AI forecasting platform available immediately
- B) Define the specific business problem, success metrics, and data availability before selecting any technology
- C) Hire a team of data scientists to build a custom model
- D) Ask the IT department to evaluate AI platforms independently

✅ **Answer: B** — Technology selection should always follow problem definition — not precede it. Leaders must first articulate: exactly what problem are we solving (stockouts? overstock? specific product categories?), how will we measure success (forecast accuracy improvement? inventory carrying cost reduction?), and what data do we have available and in what condition? Only then can an informed technology evaluation occur. Jumping to technology selection first is the leading cause of AI project failure.

What distinguishes "AI-assisted" decision-making from "AI-automated" decision-making?

- A) AI-assisted uses older models; AI-automated uses newer, more capable models
- B) AI-assisted keeps a human as the final decision-maker; AI-automated executes decisions without human approval
- C) AI-assisted is used in consumer applications; AI-automated is used only in enterprise
- D) AI-assisted requires coding knowledge; AI-automated does not

✅ **Answer: B** — This distinction is critical for governance and risk management. AI-assisted: AI provides analysis, recommendations, or options — a human decides. AI-automated: AI takes action without human approval. The appropriate model depends on the stakes, reversibility, and regulatory context of the decision. Leaders should explicitly document which decisions in their organization are AI-assisted vs. AI-automated and review this classification regularly as AI capabilities and organizational trust mature.

What is an "AI ethics board" and what is its role in an enterprise?

- A) A government-mandated regulatory body that reviews all AI deployments nationally
- B) A cross-functional internal committee that reviews high-risk AI use cases, sets ethical guidelines, and ensures AI initiatives align with organizational values
- C) An external audit firm that certifies AI systems as ethical
- D) A team of philosophy academics who write AI ethics research papers

✅ **Answer: B** — An internal AI ethics board or review committee typically includes representatives from Legal, HR, Technology, Risk, and business leadership. It reviews proposed AI use cases against ethical guidelines before deployment — particularly for high-risk applications. It establishes principles, creates review checklists, handles escalations when ethical concerns arise, and ensures accountability for AI decisions resides with identifiable humans within the organization.

Which of the following AI governance actions is MOST important for leaders in heavily regulated industries like banking or healthcare?

- A) Deploying AI as quickly as possible to gain competitive advantage before regulations tighten
- B) Documenting AI model decisions with audit trails and ensuring explainability for any AI used in regulated processes
- C) Using only open-source AI models to avoid vendor restrictions
- D) Limiting AI use to back-office functions that regulators never examine

✅ **Answer: B** — Regulated industries face heightened AI governance requirements. Regulators increasingly require organizations to explain algorithmic decisions, maintain audit trails of AI-influenced decisions, demonstrate bias testing, and prove human oversight. Leaders in these sectors must embed regulatory requirements into AI deployment processes from the start — retrofitting governance onto deployed AI is significantly more expensive and disruptive than building it in from the beginning.

The EU AI Act classifies AI systems into risk tiers. Which of the following is classified as "high-risk AI" under this framework?

- A) A spell-checker in a word processing application
- B) An AI music recommendation system on a streaming platform
- C) An AI system used to evaluate job applications or make credit scoring decisions
- D) An AI chatbot providing general cooking recipe suggestions

✅ **Answer: C** — The EU AI Act's high-risk tier includes AI used in employment (recruitment, performance evaluation), credit scoring, education, healthcare, law enforcement, immigration, and critical infrastructure. High-risk AI systems face strict requirements: technical documentation, conformity assessment, human oversight mechanisms, and registration in an EU database. Leaders must identify whether their AI deployments fall into high-risk categories and ensure compliance pathways are in place.

What does "fairness" mean in the context of responsible AI?

- A) That all users receive the same output from an AI system, regardless of context
- B) That an AI system does not systematically disadvantage individuals or groups based on protected characteristics
- C) That AI systems are available to all users at the same price
- D) That AI development teams include members from diverse backgrounds

✅ **Answer: B** — AI fairness is the property that a model's outputs do not produce unjust disparate impacts across demographic groups — race, gender, age, disability status, etc. Fairness is technically complex because there are multiple mathematically incompatible definitions (demographic parity, equalized odds, individual fairness). Leaders should require that AI vendors document which fairness definition they optimize for and how they test for it, particularly in high-stakes use cases.

What is "data minimization" in the context of AI and privacy?

- A) Reducing the size of training datasets to lower compute costs
- B) The principle of collecting and processing only the data strictly necessary for a defined AI purpose
- C) Compressing data files to reduce storage requirements
- D) Limiting the number of data sources an AI system can access

✅ **Answer: B** — Data minimization is a core GDPR principle with direct AI implications. AI systems should not collect or retain more personal data than is necessary for their stated purpose. For leaders, this means auditing what data AI systems access, challenging whether all data fields are genuinely necessary, implementing automatic data deletion schedules, and ensuring that AI systems aren't retaining personal information beyond their operational need.

What is "explainability" (XAI) and when is it most critical?

- A) The ability of AI developers to explain the source code of their models
- B) The property of an AI system that allows its decisions or recommendations to be understood and justified by humans
- C) A feature that translates AI outputs into multiple languages automatically
- D) The requirement that AI systems display confidence percentages with every output

✅ **Answer: B** — Explainability is most critical when AI influences consequential decisions affecting individuals — loan approvals, medical diagnoses, hiring decisions, criminal sentencing. In these contexts, the ability to explain why a decision was made is both an ethical obligation and increasingly a legal requirement. Techniques include LIME and SHAP (methods that identify which input features most influenced a decision). Leaders must mandate explainability requirements in AI procurement for high-stakes use cases.

A leader discovers that their AI recruitment tool has a statistically lower acceptance rate for applications from women in technical roles. What is the correct response?

- A) Conclude that fewer qualified women applied and take no action
- B) Immediately suspend the tool, conduct a bias audit, identify the root cause, remediate, and implement ongoing monitoring before redeployment
- C) Reduce the weight of gender as a variable in the model
- D) Accept the outcome as the AI is making statistically accurate predictions

✅ **Answer: B** — Disparate impact in AI hiring decisions is both an ethical failure and a legal liability under employment discrimination law in most jurisdictions. The correct response follows an incident response framework: suspend the system, investigate (was bias in training data? in feature selection? in outcome labels?), remediate the root cause (not just the symptom), test rigorously before redeployment, and implement ongoing statistical monitoring for disparate impact across demographic groups going forward.

What is "accountability" in responsible AI, and who holds it in an enterprise?

- A) Accountability means all AI errors must be publicly disclosed to regulators
- B) Accountability means clearly defined human ownership of AI systems and their outcomes — someone is responsible when AI causes harm
- C) Accountability is solely the responsibility of the AI vendor who built the model
- D) Accountability means the IT department is responsible for all AI-related incidents

✅ **Answer: B** — A core failure mode in enterprise AI governance is accountability diffusion — when AI causes harm, no individual or team is clearly responsible. Responsible AI requires explicit accountability assignment: who owns this AI system, who is responsible for monitoring its performance and fairness, and who is accountable when it fails? This must be documented, not assumed. Leaders cannot delegate accountability for AI outcomes entirely to vendors or technical teams.

What is the purpose of an "AI Impact Assessment" before deploying a new AI system?

- A) To calculate the expected return on investment of the AI deployment
- B) To systematically evaluate potential harms — to individuals, groups, and society — before an AI system goes live
- C) To assess the impact of AI on server infrastructure capacity
- D) To determine how many employees will lose their jobs due to the AI deployment

✅ **Answer: B** — An AI Impact Assessment (analogous to a Data Protection Impact Assessment under GDPR) systematically evaluates potential negative consequences of an AI deployment before launch — bias risks, privacy risks, safety risks, and societal impacts. It forces proactive thinking about harm, not reactive response after harm occurs. Leading governance frameworks (EU AI Act, NIST AI RMF) require or strongly recommend impact assessments for high-risk AI applications.

What does "transparency" mean in responsible AI practice?

- A) Publishing all source code and training data for every AI model publicly
- B) Being open about when AI is being used, how it makes decisions, and what its limitations are — with users, employees, and other stakeholders
- C) Providing real-time access to AI model weights to any requesting researcher
- D) Disclosing AI vendor contracts publicly to shareholders

✅ **Answer: B** — AI transparency operates at multiple levels: users should know when they are interacting with AI, employees should understand how AI influences decisions that affect them, customers should be told how their data is used in AI systems, and governance boards should have access to model documentation, performance metrics, and incident reports. Transparency doesn't require open-sourcing everything — it requires honest, appropriate communication with each stakeholder group.

Which of the following represents the highest risk from a data privacy perspective when deploying a third-party LLM API?

- A) The API response time being slower than expected
- B) Sensitive organizational or personal data entered into the API being used to train the vendor's model or accessed by vendor employees
- C) The model producing outputs in a different language than expected
- D) API rate limits preventing high-volume usage

✅ **Answer: B** — The default terms of consumer AI APIs often permit the provider to use inputs for model improvement, creating a direct data leakage risk when employees enter confidential, proprietary, or personal information. Enterprise contracts must explicitly prohibit training on customer data, include Data Processing Agreements (DPAs) under GDPR, and be backed by technical controls (e.g., zero-data-retention API options). This is why enterprise AI contracts require legal review before deployment.

What is "algorithmic accountability" and how does it differ from traditional IT accountability?

- A) They are identical — all software follows the same accountability standards
- B) Algorithmic accountability specifically addresses the challenge that AI systems make probabilistic decisions that can cause harm at scale, often in ways that are difficult to trace or reverse
- C) Algorithmic accountability applies only to systems built by tech companies
- D) Traditional IT accountability is more rigorous than algorithmic accountability

✅ **Answer: B** — Traditional IT systems follow deterministic rules — if something goes wrong, you can trace it to a specific line of code. AI systems are probabilistic — they make statistically likely predictions that can be wrong in systematic, hard-to-detect ways. When an AI makes biased decisions at scale, millions of individuals can be affected before the problem is detected. This scale and opacity require specific accountability frameworks that traditional IT governance doesn't address.

What is the most important thing a leader should do immediately if an AI system at their organization causes a public harm incident?

- A) Delete all records of the AI system to limit liability
- B) Issue a public statement blaming the AI vendor
- C) Take the system offline or restrict its use, assign an incident owner, investigate root cause, communicate transparently, and remediate before redeployment
- D) Wait for regulatory guidance before taking any action

✅ **Answer: C** — AI incident response follows a clear playbook: immediate containment (restrict or suspend the system), clear ownership assignment (one person leads the response), transparent communication (internal and external, appropriate to the severity), root cause investigation (not surface-level blame), technical and process remediation, and documented learning that updates governance processes. Leaders who respond to AI incidents with cover-up behaviors face significantly worse regulatory and reputational consequences than those who respond with transparency and swift action.

Under the EU AI Act, what is classified as "unacceptable risk" AI that is completely prohibited?

- A) AI used in social media content recommendation
- B) AI-powered customer service chatbots
- C) AI systems that manipulate people through subliminal techniques, enable real-time mass biometric surveillance in public spaces, or implement social scoring by governments
- D) AI tools used for employee productivity monitoring

✅ **Answer: C** — The EU AI Act's prohibited category includes: AI that exploits subconscious vulnerabilities to manipulate behavior, real-time remote biometric identification in public spaces (with narrow exceptions), AI-based social scoring systems by public authorities, and systems that exploit children or vulnerable groups. These are considered so fundamentally harmful that no business justification is acceptable. Leaders in EU markets must ensure no organizational AI deployments fall into these categories.

What is "consent" in AI data use, and why is it challenging to obtain properly?

- A) A one-time agreement that permits all future uses of data once obtained
- B) Informed, specific, freely given agreement by individuals to a clearly defined use of their data — challenging because AI use cases evolve and data is often repurposed beyond its original consent scope
- C) A legal contract between the AI vendor and the organization deploying the AI
- D) Automated acceptance of terms of service during account creation

✅ **Answer: B** — Genuine AI consent requires individuals to understand specifically how their data will be used — a challenge because AI systems often find new uses for data that were not anticipated when consent was obtained. Re-purposing data collected for one AI use case for a different AI application without new consent is a common compliance failure. Leaders must implement data governance that tracks consent scope and prevents unauthorized repurposing, especially as AI capabilities and use cases evolve.

What is "AI safety" in the enterprise context — and how does it differ from AI ethics?

- A) They are the same concept with different names
- B) AI ethics addresses fairness and values alignment; AI safety focuses on preventing AI systems from causing harm through failures, unexpected behaviors, or misuse — including both technical and operational safeguards
- C) AI safety refers only to cybersecurity protections around AI systems
- D) AI ethics applies to consumer AI; AI safety applies only to military AI

✅ **Answer: B** — AI ethics addresses normative questions: is this AI doing the right thing? AI safety addresses operational questions: is this AI working as intended and failing safely? Enterprise AI safety includes: robust testing before deployment, fail-safes that revert to human processes when AI confidence is low, monitoring for unexpected outputs, adversarial testing (what happens if someone tries to manipulate this system?), and clear procedures for safe shutdown. Both ethics and safety are essential and complementary components of responsible AI.

A leader in a financial services firm is evaluating an AI credit scoring model from a vendor. Which is the MOST important due diligence question regarding responsible AI?

- A) "How fast does the model process applications?"
- B) "Which famous clients use your model?"
- C) "How has the model been tested for disparate impact across racial and demographic groups, and what are the documented results?"
- D) "What is your company's revenue and how long have you been in business?"

✅ **Answer: C** — Credit scoring AI in financial services has a documented history of perpetuating or amplifying racial and socioeconomic disparities. Due diligence must include: documented fairness testing across protected classes, the statistical methodology used (which fairness definition), the specific test results and pass/fail thresholds, and ongoing monitoring commitments. This is not just ethical due diligence — it is regulatory due diligence under laws like ECOA (Equal Credit Opportunity Act) in the US and equivalent legislation globally.

GSDC
Global Skill Development Council

CERTIFIED GENERATIVE AI FOR LEADERS

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING OBJECTIVE

- Gain insights into autonomous decision-making processes
- Apply knowledge using ready-to-implement templates
- Demonstrate ability to work with Agentic AI models
- Validate your skills wit

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org