PRACTICE EXAM 6: AIGP SIMULATION (100 QUESTIONS)

1. A governance professional joins a new organization and discovers the company has 30 deployed AI systems but no AI inventory. The professional's manager asks them to "get governance started." Which of the following tasks should be the professional's VERY FIRST action?

A. Draft a comprehensive AI governance policy manual that establishes responsible AI principles, governance roles, and compliance procedures for the entire organization

B. Schedule meetings with each business unit to discuss responsible AI principles and begin building a governance culture across the organization

C. Conduct an inventory of all 30 AI systems — documenting what each does, who it affects, what data it processes, and a preliminary risk estimate — because governance cannot be designed or prioritized without understanding what needs to be governed

D. Select and implement an AI governance software platform that can automate monitoring, documentation, and compliance tracking across all deployed systems

2. A governance professional reviews a vendor's model card for a high-risk AI hiring tool. The model card states the system was "validated on a diverse dataset." The professional requests the specific demographic breakdown of the validation dataset. The vendor responds that providing demographic breakdowns would "violate data protection principles." How should the professional evaluate this response?

A. The vendor's response is likely an evasion — aggregate demographic statistics about a validation dataset (e.g., percentage by gender, age range, ethnicity) do not constitute personal data and can be shared without violating data protection principles, and the deployer needs this information to fulfill its independent governance obligations

B. The vendor's response is correct because all demographic information about datasets is classified as personal data under GDPR regardless of whether individual records are identifiable

C. The vendor's response is acceptable because the statement "diverse dataset" provides sufficient assurance for the deployer to proceed with deployment

D. The professional should accept the vendor's position and rely on the deployer's own post-deployment monitoring to evaluate the system's fairness

3. A governance professional discovers that the organization's AI-powered email filtering system — originally deployed to filter spam — has been modified by the IT team to also scan employee emails for "negative sentiment about the company." This modification was made without governance review. The IT team argues the system already had email access, so no additional governance was needed. What is the professional's MOST accurate assessment?

A. The IT team is correct because the system already had authorized access to email content and the modification does not change the system's data inputs

B. The modification only requires a documentation update to reflect the expanded functionality without any additional governance review

C. The modification requires only notification to the data protection officer because the change involves processing personal data for a new purpose

D. The modification constitutes a secondary use requiring comprehensive governance review — the system's purpose has changed from spam filtering to employee surveillance, triggering new privacy, employment law, and potentially EU AI Act considerations

4. During a governance committee meeting, a product manager presents a proposal for an AI system that would analyze patients' social media posts to predict medication non-adherence. The system would alert healthcare providers when patients appear likely to stop taking prescribed medications. Two committee members support the proposal for its patient safety benefits. Two oppose it for privacy reasons. The governance professional must advise the committee. What approach BEST reflects governance principles?

A. Side with the supporters because patient safety is a responsible AI principle that automatically overrides privacy concerns in healthcare contexts

B. Frame the decision as requiring proportionality analysis — evaluating whether the patient safety benefit justifies the privacy intrusion, whether less invasive alternatives exist, whether patients can meaningfully consent, and whether the system's predictions are accurate enough to warrant acting on them

C. Side with the opponents because processing social media data for healthcare purposes is prohibited under all data protection frameworks

D. Defer the decision to the healthcare provider's medical ethics board because AI governance committees lack the clinical expertise to evaluate patient safety implications

5. A governance professional is conducting a quarterly review of a deployed AI loan approval system. The monitoring dashboard shows all metrics within acceptable ranges. However, the professional notices that the volume of applications has increased by 300% since deployment, and the system is now processing applications from a market segment (gig economy workers) that was not represented in the original training data. Should the professional flag this as a governance concern despite the stable metrics?

A. Yes — the system is now processing a population not represented in its training data, which means the stable aggregate metrics may mask poor performance or unfair outcomes for gig economy workers specifically, and disaggregated analysis for this new segment is needed

B. No, because the monitoring metrics are within acceptable ranges and the system has demonstrated it can handle increased volume without degradation

C. Yes, but only because the 300% volume increase may exceed the system's computational capacity and create latency issues affecting response times

D. No, because gig economy workers are not a protected class under nondiscrimination law and their inclusion does not create a fairness concern

6. A governance professional receives an urgent request from the CEO to deploy an AI customer service system "by end of week" to handle a sudden surge in support volume. The system has completed development and basic testing but has not undergone impact assessment, bias testing, or governance committee review. The CEO argues that customer satisfaction is at risk. What is the MOST appropriate governance response?

A. Deploy the system as requested because the CEO has ultimate decision-making authority and governance must not impede executive business decisions

B. Refuse to deploy under any circumstances because skipping governance review for a customer-facing system violates the organization's governance policy regardless of business pressure

C. Conduct a rapid risk assessment, identify the minimum essential governance controls that can be completed within the timeframe, deploy with those controls plus enhanced monitoring and clear rollback capability, and schedule the full governance review for completion within 30 days

D. Deploy the system but limit it to a non-customer-facing internal pilot for the first month while the full governance review is completed in parallel

7. A governance professional discovers that three different teams within the organization have independently purchased the same AI analytics tool from the same vendor — each with different contractual terms, different data processing agreements, and different governance oversight levels. What governance structural issue does this reveal?

A. A training deficiency, because the three teams were not adequately trained on the organization's preferred vendor list and procurement procedures

B. A monitoring failure, because the organization's AI monitoring system should have detected duplicate vendor relationships automatically

C. A documentation failure, because each team should have registered their purchase in the organization's software asset management system

D. The absence of a centralized AI procurement process and AI inventory — without organizational visibility into AI acquisitions, governance cannot ensure consistent vendor assessment, contractual protections, data governance, and compliance standards across the organization

8. A governance professional is reviewing a DPIA for an AI system that analyzes employee biometric data (fingerprints) for building access control. The DPIA was completed by the IT security team and concludes that the system poses "low privacy risk because fingerprints are used only for access control, not identification." The professional identifies a significant error. What is it?

A. The DPIA incorrectly assumed that biometric access control is classified as low risk, when the EU AI Act classifies all biometric systems as prohibited

B. The DPIA's conclusion is incorrect because using fingerprints for access control IS identification — biometric access control verifies identity by matching a fingerprint against enrolled records, constituting processing of special category data under GDPR Article 9 regardless of the stated purpose

C. The DPIA was conducted by the wrong team — IT security teams are prohibited from conducting DPIAs under GDPR Article 35

D. The DPIA's error is that it failed to include a cost-benefit analysis comparing biometric access control against alternative technologies

9. A governance professional is asked to review a contract for an AI vendor's cloud-based sentiment analysis service. The professional notes that the contract does not address what happens to the organization's data if the vendor is acquired by another company. Why is this omission a governance concern?

A. The omission is not a governance concern because acquisitions are regulated by antitrust law and do not affect data processing agreements

B. The omission only matters if the acquiring company is headquartered outside the European Economic Area, because intra-EEA acquisitions do not change data governance obligations

C. The omission is a governance concern because a change of ownership could expose the organization's data to a competitor, change the vendor's data processing practices, or transfer the data to an entity with different privacy standards — and without contractual protections, the organization has no remedy

D. The omission is only relevant if the organization processes special category data through the vendor's service

10. A governance professional is evaluating an AI system that uses facial analysis to estimate customers' ages at self-checkout registers to verify eligibility for age-restricted purchases (alcohol, tobacco). The system does not store facial images — it processes them in real time and discards them. The development team argues this means no personal data is processed. Is this argument correct?

A. No — the real-time processing of facial images to estimate age constitutes personal data processing under GDPR even if the images are not stored, because the processing involves identifying or inferring characteristics of a natural person from their biometric features

B. Yes, because GDPR only applies to personal data that is stored in a filing system, and real-time processing without storage does not constitute regulated processing

C. Yes, because age estimation is not the same as identification, and only facial recognition systems that identify specific individuals constitute personal data processing

D. No, but only because the system processes data of minors who cannot provide valid consent, not because of the facial processing itself

11. A governance professional is reviewing the organization's AI incident response plan. The plan assigns the head of the data science team as the sole decision-maker for incident severity classification and response actions. The professional identifies a governance weakness. What is it?

A. The head of data science should not be involved in incident response at all because technical teams have conflicts of interest when responding to incidents involving their own systems

B. The severity classification should be determined by the AI vendor rather than by any internal team member because the vendor has the deepest technical knowledge

C. Incident response plans should assign the CEO as the sole decision-maker to ensure executive-level accountability for all AI-related incidents

D. The data science team leader has an inherent conflict of interest as the person responsible for the systems that may have caused the incident — incident response decisions should involve cross-functional authority including governance, legal, and business stakeholders

12. A governance professional is advising the organization on how to handle a situation where an AI system's monitoring reveals a gradual increase in data drift over three consecutive months but no corresponding change in performance or fairness metrics. The operations team wants to ignore the drift because "nothing is broken." What should the professional recommend?

A. Agree with the operations team because data drift without performance impact indicates the model is robust and does not require intervention

B. Investigate the cause of the drift, increase monitoring frequency, and prepare retraining plans — because data drift is a leading indicator that often precedes performance degradation, and the fact that impact has not yet materialized does not mean it will not materialize

C. Immediately retrain the model because any data drift above detection thresholds requires mandatory model replacement under the EU AI Act

D. Report the data drift to the national competent authority as a potential serious incident because systematic distribution changes may indicate adversarial data manipulation

13. A governance professional discovers that an AI vendor has been sending weekly marketing emails to the organization's employees who interact with the vendor's AI system, using contact information obtained through the system's user authentication process. The vendor agreement does not address the vendor's use of employee contact information for marketing. What governance issue does this raise?

A. The vendor is likely using employee contact data for a purpose (marketing) beyond what was authorized for the AI system's operation (user authentication), potentially violating data protection principles and constituting unauthorized processing of employee personal data

B. The vendor's marketing emails are a standard business practice and do not raise governance concerns as long as employees can unsubscribe from the mailing list

C. The governance issue is limited to the IT department's failure to filter the vendor's marketing emails through the organization's email security system

D. The vendor's use of employee contact information is governed exclusively by anti-spam regulations and falls outside the scope of AI governance

14. A governance professional is reviewing the testing protocol for an AI system designed to detect fraudulent insurance claims. The protocol includes performance testing on a held-out test set and fairness testing across gender groups. The professional identifies a gap. What is the MOST significant missing element?

A. The protocol should include user acceptance testing with insurance adjusters to verify that the system's interface is intuitive and easy to use

B. The protocol should include load testing to verify the system can handle peak claim volumes during natural disaster seasons

C. The protocol should include intersectional fairness testing — evaluating outcomes across combinations of protected characteristics (gender × race, gender × age, race × income level) — because single-axis testing can miss disparities that emerge at intersections

D. The protocol should include a comparison against the industry benchmark fraud detection rate to verify competitive positioning

15. A governance committee is reviewing a proposal to deploy an AI system that would analyze job candidates' voice patterns during phone interviews to assess "confidence," "trustworthiness," and

"leadership potential." The vendor claims the system has been "scientifically validated." What should the governance committee's FIRST concern be?

A. Whether the system's voice analysis discriminates against candidates with speech impediments, accents, or voice characteristics associated with specific demographic groups

B. Whether the system complies with wiretapping laws that may require consent from all parties before recording and analyzing phone conversations

C. Whether the vendor's pricing model is competitive with alternative interview assessment tools available in the market

D. Whether the constructs being measured — confidence, trustworthiness, and leadership potential — can be validly assessed from voice patterns at all, because deploying a system that measures scientifically invalid constructs creates harm regardless of technical accuracy

16. A governance professional is conducting a post-incident review after an AI content moderation system incorrectly removed thousands of posts documenting human rights abuses, classifying them as "violent content." The development team has already fixed the classification error and redeployed the system. The professional argues that the technical fix alone is an incomplete response. Why?

A. The professional is incorrect — once the classification error is fixed, the incident is resolved and no further action is required

B. The post-incident review should also evaluate why the monitoring system did not detect the mass removal of legitimate content, whether affected users should be notified and their content restored, whether the incident reveals systemic gaps in how the content moderation system handles documentation of real-world violence versus violent content, and what governance improvements would prevent recurrence

C. The only additional step needed is to update the system's model card to document the classification error as a known limitation

D. The professional's concern is valid only because the removed content related to human rights — a similar mass removal of other content types would not require post-incident governance review

17. A governance professional is evaluating two approaches to human oversight for an AI medical triage system. Approach A has a nurse review every AI triage recommendation before it is acted upon.

Approach B has the AI system act on its recommendations immediately, with a physician reviewing a statistical sample of decisions daily. For a system that assigns urgency levels in an emergency department where minutes matter, which approach is MOST governance-appropriate?

A. Neither approach is ideal in isolation — the governance framework should implement risk-proportionate oversight: Approach B (immediate AI action with statistical review) for lower-acuity classifications where delays are less harmful, and Approach A (human review before action) for high-acuity classifications where the consequences of error are severe

B. Approach A, because every medical decision must be reviewed by a human before any action is taken, regardless of the clinical urgency or the time constraints

C. Approach B, because statistical sampling provides more comprehensive oversight than individual case review and is the EU AI Act's required approach for medical AI

D. A hybrid approach where the AI system makes all decisions autonomously and a weekly management report summarizes the system's performance for administrative review

18. A governance professional is reviewing the organization's data retention policy for AI training data. The current policy retains all training data indefinitely "in case it is needed for model retraining." The professional identifies this as a governance issue. On what basis?

A. Indefinite retention is required by the EU AI Act for all training data associated with high-risk AI systems and the professional's concern is unfounded

B. The professional's concern is limited to storage costs, which are not a governance issue but an operational budget consideration

C. Indefinite retention of training data may violate data minimization and storage limitation principles under GDPR — personal data should be retained only for as long as necessary for the specified purpose, and "in case it might be needed" is not a sufficiently defined purpose to justify indefinite retention

D. Indefinite retention is only a governance concern if the training data contains biometric or health data

19. A governance professional is advising the organization on an AI system that generates synthetic voices for audiobook production. The system was trained on recordings of professional voice actors. The organization's legal team has reviewed the licensing agreements with the voice actors and confirmed that the agreements authorize the use of recordings for "AI research and development." The organization

now wants to use the trained system commercially to produce audiobooks without paying royalties to the original voice actors. What governance issue should the professional raise?

A. No governance issue exists because the legal team has confirmed that the licensing agreements authorize AI use and commercial deployment falls within the scope of "research and development"

B. The governance issue is limited to ensuring that the AI-generated audiobooks are clearly labeled as synthetic content under the EU AI Act's transparency provisions

C. The governance issue is only relevant if any of the voice actors are EU citizens, because non-EU voice actors are not protected by data protection or IP regulations

D. Commercial audiobook production likely exceeds the "AI research and development" purpose authorized in the licensing agreements — the governance professional should raise the purpose limitation concern and recommend obtaining explicit authorization for commercial use before deploying the system

20. A governance professional receives a report from the monitoring team showing that an AI customer segmentation system's performance has been stable for 18 months. The monitoring tracks overall accuracy, precision, and recall. The professional requests additional monitoring for fairness metrics disaggregated by customer demographics. The monitoring team pushes back, arguing that adding fairness metrics is unnecessary because the system does not make decisions about individuals — it only groups customers for marketing purposes. How should the professional respond?

A. The professional should insist on fairness monitoring because customer segmentation based on AI can produce discriminatory outcomes — if the system systematically places customers from certain demographic groups into less favorable segments, those customers may receive inferior service, fewer offers, or different pricing, constituting discriminatory treatment even though no individual "decision" is made

B. The monitoring team is correct because fairness metrics are only required for AI systems that make consequential decisions about individuals, and customer segmentation does not meet this threshold

C. The professional should defer to the monitoring team's technical expertise because they have deeper knowledge of what metrics are appropriate for segmentation systems

D. The professional should agree to defer fairness monitoring for six months and revisit the question if customer complaints about discrimination arise during that period

21. A governance professional is reviewing an organization's AI governance budget allocation. The budget dedicates 80% to pre-deployment governance activities (impact assessments, testing, documentation, governance committee reviews) and 20% to post-deployment activities (monitoring, maintenance, incident response, periodic reassessment). The professional recommends rebalancing. What is the governance rationale for this recommendation?

A. The budget should allocate 80% to post-deployment activities because most governance costs occur after deployment, and the current allocation underinvests in the continuous monitoring, maintenance, and reassessment that constitute the longest and most consequential phase of the AI lifecycle

B. Post-deployment activities typically require greater sustained investment than pre-deployment activities because they continue for the entire operational life of the system — monitoring, maintenance, incident response, and periodic reassessment need ongoing funding rather than one-time investment

C. The budget should be equally split 50/50 between pre-deployment and post-deployment activities because the EU AI Act requires equal investment in both phases

D. The budget allocation is appropriate because pre-deployment governance is more important than post-deployment governance and should receive the majority of funding

22. A governance professional is asked to evaluate whether an AI system's training data was collected ethically. The data consists of social media posts scraped from public profiles. The professional identifies several governance concerns. Which concern is MOST fundamental?

A. The social media posts may contain copyrighted content such as original creative writing, photographs, and artwork that require licensing for AI training

B. The social media platform's terms of service may prohibit scraping user content for commercial AI training, creating a contractual violation regardless of the data's public availability

C. The posts may contain personal data of individuals who did not consent to their content being used for AI training, and public availability does not constitute consent under GDPR

D. The public availability of the data does not resolve the purpose limitation question — individuals posted on social media for social communication purposes, and using that content for AI training is a fundamentally different purpose that may not be compatible with the original purpose, regardless of the data's public visibility

23. A governance professional is advising an organization that wants to deploy the same AI hiring system in both France and Saudi Arabia. The system evaluates candidates' qualifications and produces a ranked shortlist. The professional identifies that the concept of "protected characteristics" differs between jurisdictions. What governance implication does this create?

A. The organization should use the same fairness metrics in both jurisdictions because mathematical definitions of fairness are universal and do not vary by jurisdiction

B. The organization should deploy the system only in the jurisdiction with fewer protected characteristics to minimize governance complexity and compliance costs

C. Bias testing must be calibrated to each jurisdiction's protected characteristics, cultural context, and legal requirements — the groups against which discrimination must be evaluated differ between France and Saudi Arabia, and fairness testing valid in one context may not address the relevant fairness dimensions in the other

D. The organization should defer deployment in both jurisdictions until international standards organizations establish a universal set of protected characteristics

24. A governance professional discovers that the organization's deployed AI system for automated invoice processing has been approving invoices that contain mathematical errors — the system verifies vendor identity and purchase order matching but does not check whether the invoice amounts are arithmetically correct. Over the past year, approximately €400,000 in overpayments have occurred. The system was never designed to check arithmetic. What type of governance failure is this?

A. A gap in the use case assessment and system design — the system's intended purpose (automated invoice processing) was defined too narrowly, omitting a fundamental requirement (arithmetic verification) that any invoice processing system should include, resulting in a design that automates approval without basic quality checks

B. A monitoring failure because the post-deployment monitoring system should have detected the pattern of overpayments through anomaly detection

C. A vendor liability issue because the AI vendor should have included arithmetic verification as a standard feature in its invoice processing product

D. A training data failure because the training data did not contain sufficient examples of invoices with mathematical errors for the system to learn to detect them

25. A governance professional is reviewing the organization's approach to AI model versioning. The organization maintains only the current production version of each AI model — when a model is retrained or updated, the previous version is deleted. The professional identifies this as a governance risk. Why?

A. The organization is required by the EU AI Act to maintain all historical versions of every AI model permanently in an archived format accessible to regulatory authorities

B. Without previous versions, the organization cannot conduct comparative analysis when issues arise (did the problem exist in the previous version?), cannot roll back to a known-good version if an update causes problems, and cannot support incident investigation that requires understanding how the system's behavior changed over time

C. Maintaining only the current version violates ISO/IEC 42001's requirement that organizations preserve every model version for the duration of the system's operational life plus five years

D. The risk is limited to intellectual property concerns because previous model versions may contain proprietary innovations that competitors could exploit if the current version is compromised

26. A governance professional is reviewing a proposal to use an AI system for predictive maintenance in a nuclear power plant. The system would analyze sensor data to predict equipment failures. The professional notes that the proposal's risk assessment classifies the system as "moderate risk" based on the organization's standard risk classification framework. What is the professional's MOST likely concern?

A. The moderate risk classification is appropriate because predictive maintenance systems do not directly control safety-critical equipment and only provide advisory recommendations

B. The professional's concern is limited to ensuring that the AI vendor has the appropriate security clearances required for operating in nuclear facility environments

C. The professional's concern is limited to ensuring that the system's predictions are reviewed by qualified nuclear engineers before any maintenance decisions are made

D. The standard risk classification framework may not adequately account for the cascading consequences of AI errors in a nuclear context — a missed failure prediction or a false negative in a nuclear power plant carries consequences of an entirely different magnitude than in a standard industrial setting, potentially warranting a higher risk classification

27. A governance professional is advising an organization that has developed a proprietary AI model and wants to release it as open source. The organization believes this will demonstrate transparency and build trust. The professional identifies several governance considerations. Which is MOST critical to address BEFORE the release?

A. Once open-sourced, the organization permanently loses trade secret protection for the model and cannot control downstream uses — including potentially harmful applications such as surveillance, discrimination, or weapons development — and must evaluate whether the transparency benefits outweigh these irreversible consequences

B. Open-sourcing requires the organization to provide perpetual free technical support to all users of the model under international open-source licensing obligations

C. The organization must obtain approval from the European AI Office before releasing any AI model as open source within the EU market

D. Open-source release is only a governance concern if the model was classified as high-risk under the EU AI Act

28. A governance professional is conducting an annual review of the organization's AI governance training program. Completion rates are 95% across the organization. However, when the professional interviews employees, they find that most cannot describe the organization's AI governance policies, cannot identify which AI systems they interact with, and do not know how to report AI-related concerns. What does this finding reveal?

A. The training program is effective because 95% completion demonstrates comprehensive organizational coverage, and employee retention of specific policies is not a training program's responsibility

B. The finding reveals only that the training content needs to be updated with more engaging multimedia elements to improve knowledge retention

C. The finding reveals that the training program measures completion rather than effectiveness — high completion rates with low comprehension and behavioral change indicate the training is not achieving its governance purpose, and the program needs redesign focused on practical competence, not just completion metrics

D. The finding is expected because AI governance is a specialized topic and frontline employees should not be expected to understand governance policies in detail

29. A governance professional is reviewing an AI system that automatically generates employment references for former employees based on their HR records. The system produces standardized references that include performance ratings, disciplinary history, and attendance records. The professional identifies multiple governance concerns. What is the MOST significant concern?

A. The system may violate data protection principles by processing more personal data than necessary for a reference — not all HR data (disciplinary records, medical absences) is appropriate for external disclosure

B. The system generates references that are not sufficiently personalized because they rely on structured data rather than qualitative assessments from supervisors

C. The system should only generate positive references because including any negative information in an automated reference could constitute defamation

D. The system violates GDPR's automated decision-making provisions because generating an employment reference constitutes a solely automated decision with significant effects on the former employee's future employment prospects

30. A governance professional discovers that the organization's AI fraud detection system has been operating for eight months with a logging infrastructure failure — no decision logs exist for approximately 200,000 fraud determinations. The system has been flagging transactions and some customers have had accounts frozen based on AI-generated fraud alerts. What is the professional's MOST immediate governance concern?

A. The organization's IT infrastructure team should be disciplined for allowing the logging failure to persist for eight months without detection

B. The organization cannot demonstrate the basis for any of the 200,000 decisions, cannot verify whether those decisions were fair or accurate, cannot respond to customer complaints or regulatory inquiries about specific decisions, and may have frozen customer accounts based on unauditable AI determinations — creating significant compliance, accountability, and potential liability exposure

C. The logging failure is a minor technical issue that can be resolved by re-enabling logging and documenting the gap in the system's maintenance records

D. The professional's only concern should be reporting the logging failure to the national competent authority as a serious incident under the EU AI Act

31. A governance professional is advising on an AI system that generates synthetic data for training a healthcare AI model. The synthetic data is created from real patient records and is designed to preserve the statistical properties of the original data while removing identifiable information. The development team argues that since the synthetic data is not real patient data, no data governance controls are needed. What is the professional's assessment?

A. Synthetic data generated from real patient records requires governance oversight — the generation process may not fully remove identifiable information, the synthetic data may closely resemble specific patients, and the fidelity of synthetic data to the original population must be validated to ensure the downstream model does not learn artifacts rather than genuine patterns

B. The development team is correct because synthetic data by definition does not contain personal data and falls entirely outside the scope of data governance

C. Synthetic data requires governance only if the generation model itself is classified as high-risk under the EU AI Act

D. The professional's only concern should be verifying that the synthetic data generation algorithm is not patented by another organization

32. A governance professional is reviewing a deployed AI chatbot that provides customer support for a telecommunications company. The chatbot has been trained on historical customer service transcripts. The professional discovers that the chatbot occasionally tells customers they are eligible for discounts or plan changes that the company does not actually offer — because similar offers appeared in historical transcripts from promotional periods that have ended. What type of governance issue is this?

A. A cybersecurity vulnerability because the chatbot's incorrect discount offers could be exploited by customers to fraudulently obtain unauthorized discounts

B. A model architecture issue that requires the chatbot to be redesigned using a retrieval augmented generation approach that grounds responses in current offer databases

C. A training data governance issue — the historical transcripts contain outdated information that the chatbot presents as current, and the system lacks mechanisms to distinguish between current and expired offers, creating false representations to customers that may constitute deceptive practices under consumer protection law

D. An acceptable limitation of generative AI systems that can be addressed by adding a disclaimer informing customers that the chatbot's information may not reflect current offers

33. A governance professional is advising the organization's legal team on a patent application for an AI invention. The AI system independently discovered a novel chemical compound with pharmaceutical potential during automated molecular analysis. No human scientist directed the AI to search for this specific compound or contributed creative input to the discovery. The legal team wants to list the AI system as co-inventor. What should the professional advise?

A. The AI system should be listed as sole inventor because it made the discovery without human creative contribution

B. Patent offices in most jurisdictions require a natural person as the inventor — the organization should identify the human researchers who designed the AI system, defined its research parameters, or interpreted its results as the inventors, and document the human contribution to support the patent application

C. The invention cannot be patented under any circumstances because AI-discovered compounds are categorically excluded from patent protection worldwide

D. The AI system should be listed as co-inventor alongside the organization's CEO because executive leadership authorized the research program

34. A governance professional is evaluating the organization's AI vendor management practices. The organization currently assesses vendors at the time of procurement but does not reassess them during the contract period. A critical AI vendor recently experienced a data breach affecting multiple clients. The organization was not notified for three weeks because the vendor agreement did not include breach notification terms. What governance improvement is MOST urgently needed?

A. The organization should maintain a larger roster of backup vendors so that it can immediately switch to an alternative provider when a vendor experiences a data breach

B. The organization should require all AI vendors to carry cybersecurity insurance policies with coverage limits specified in the vendor agreement

C. The organization should implement mandatory annual security audits of all AI vendor facilities conducted by the organization's own security team

D. AI vendor agreements must include specific breach notification timeframes, and the organization must implement ongoing vendor monitoring — not just procurement-time assessment — to detect changes in vendor security posture, financial stability, and governance practices during the contract period

35. A governance professional is conducting a tabletop exercise for AI incident response. The scenario involves the discovery that an AI system used for credit scoring has been producing systematically biased outcomes against a protected group for four months. During the exercise, the team debates whether to suspend the system immediately or continue operating while a fix is developed. What governance principle should guide this decision?

A. The decision should be guided by the containment principle — when a system is producing ongoing harm to a protected group, the default response is to stop the harm (suspend the system or implement manual review) while remediation is developed, rather than allowing harm to continue during the remediation period

B. The system should continue operating because suspending it would create customer service disruption that affects all users, not just the affected protected group

C. The decision should be deferred to the AI vendor because the vendor has the technical expertise to determine whether the bias can be fixed without suspending the system

D. The system should continue operating with a public disclosure that bias has been identified, allowing affected customers to opt out of AI-based credit scoring

36. A governance professional is reviewing an AI system that recommends personalized learning paths for employees in a corporate training program. The system analyzes employees' past performance reviews, skill assessments, and career aspirations to recommend courses. The professional discovers that the system's recommendations are strongly influenced by employees' current job titles — employees in junior roles are primarily recommended technical skills training, while employees in senior roles are recommended leadership and strategy training. An employee in a junior role who has expressed interest in leadership development is consistently directed away from leadership courses. What governance issue does this pattern raise?

A. No governance issue exists because the system is correctly matching training recommendations to employees' current roles and responsibilities

B. The system creates a ceiling effect that reinforces existing organizational hierarchies by directing development resources based on current position rather than individual aspirations and potential

C. The governance issue is limited to verifying that the recommendation system's training data accurately reflects the organization's current course catalog and enrollment criteria

D. The pattern is only a governance concern if the system's recommendations are mandatory — if employees can override the recommendations and enroll in any course, no governance intervention is needed

37. A governance professional is evaluating a vendor proposal for an AI system that would monitor call center employees' emotional states during customer calls. The system would analyze vocal tone, speech patterns, and word choice to generate real-time "emotional wellness scores" displayed on supervisors' dashboards. The vendor markets the system as a "workplace wellness tool." Under the EU AI Act, what regulatory classification concern should the professional raise?

A. The system is classified as minimal risk because it monitors employees for wellness purposes rather than performance evaluation

B. The system is classified as limited risk, requiring only that employees be informed that emotion analysis is being conducted during their calls

C. The system falls within a standard call quality monitoring framework and does not constitute an AI system requiring classification under the EU AI Act

D. The EU AI Act restricts emotion recognition in workplace settings — regardless of how the vendor markets the system, analyzing employees' emotional states during work activities falls within the Act's workplace emotion recognition provisions

38. A governance professional is advising on the design of a human oversight mechanism for an AI system that processes asylum applications. The system produces risk scores that influence how quickly applications are processed. The governance committee has proposed that a single immigration officer review all AI-generated risk scores. The professional identifies a problem with this design. What is it?

A. The immigration officer should be replaced with an AI review system that can evaluate the primary AI system's outputs more consistently and at greater scale

B. A single reviewer processing all AI-generated scores will experience reviewer fatigue, may develop over-reliance on the AI outputs over time (automation bias), and creates a single point of failure — the oversight mechanism should include multiple reviewers, rotation schedules, independent spot-checking, and metrics tracking whether the reviewer is genuinely evaluating scores or rubber-stamping them

C. The oversight mechanism should only review scores below a certain threshold because high-risk scores can be trusted to be accurate without human review

D. The oversight mechanism is adequate as long as the immigration officer has completed the required AI literacy training program

39. A governance professional is reviewing the organization's AI system documentation. The professional discovers that model cards for three high-risk systems were last updated when the systems were originally deployed — two years ago for one system, 18 months ago for another, and 14 months ago for the third. All three systems have undergone at least one retraining cycle since their model cards were created. What governance requirement has been violated?

A. Model cards must be updated whenever significant changes occur — including model retraining, data source changes, and performance shifts — because a model card that describes the original system no longer accurately represents the current system after retraining, making it inadequate for compliance, deployer communication, and audit purposes

B. Model cards are one-time documents created at deployment and are not required to be updated after the initial release

C. Model cards only need to be updated if the AI system's risk classification changes from one tier to another under the EU AI Act

D. Model card updates are the exclusive responsibility of the AI vendor and the deploying organization has no obligation to ensure model cards remain current

40. A governance professional is conducting a gap analysis between the organization's current AI governance practices and ISO/IEC 42001 requirements. The organization has strong technical AI practices but no formal management review process. What does the absence of management review mean for the governance program?

A. The absence is immaterial because management review is a formality that adds no practical governance value beyond what strong technical practices already provide

B. Management review is only required for organizations with more than 500 employees and the organization falls below this threshold

C. Without management review, senior leadership does not systematically evaluate whether the AI governance program is achieving its objectives, whether resources are adequate, whether policies need updating, and whether the organization's AI risk posture is acceptable — meaning strategic governance decisions are not being made

D. The gap can be resolved by having the data science team present quarterly technical performance reports to the CTO, which satisfies ISO 42001's management review requirement

41. A governance professional is advising an organization that wants to use customer complaint data to train an AI system that will predict which products are likely to generate future complaints. The complaint data includes customers' names, contact information, detailed descriptions of their issues, and emotional language expressing frustration. The data science team plans to use the full complaint text as training data. What data governance concern should the professional raise FIRST?

A. The complaint data should be preprocessed to remove personal identifiers before training, and the emotional language may need to be evaluated for potential bias impact — but the FIRST concern is whether the organization's privacy notice authorized the use of complaint data for AI training purposes

B. The professional should ensure the complaint data is complete and contains no missing values before it is used for model training

C. The emotional language in complaints should be removed because it will skew the model's predictions toward products that generate the most emotionally intense complaints rather than the most frequent complaints

D. The professional should ensure that the AI system's predictions are not shared with the product development team because sharing complaint-derived insights across departments violates data protection principles

42. A governance professional is reviewing the organization's approach to AI deactivation. The organization's policy states: "AI systems will be deactivated when they are no longer needed." The professional identifies this policy as inadequate. What SPECIFIC deficiency does the professional's assessment address?

A. The policy should specify a maximum operational lifespan for all AI systems (e.g., five years) after which they must be automatically deactivated regardless of performance

B. The policy should require that deactivated AI systems be donated to academic institutions for research purposes rather than being deleted

C. The policy is adequate because "no longer needed" is a sufficiently clear criterion for initiating deactivation and no additional specificity is required

D. The policy lacks specific deactivation criteria (performance degradation thresholds, regulatory non-compliance triggers, technology obsolescence indicators), deactivation procedures (stakeholder notification, transition planning, data disposition), and post-deactivation requirements (documentation archival, lessons-learned review) — making it unimplementable in practice

43. A governance professional receives a request from a researcher at a university who wants access to the organization's AI system's training data for an independent bias study. The researcher plans to publish the findings publicly. The organization's AI vendor has told the organization that sharing training data is prohibited under the vendor agreement's confidentiality provisions. How should the professional navigate this situation?

A. Evaluate whether the vendor agreement actually prohibits sharing aggregate statistics or anonymized subsets of training data (as opposed to the complete raw dataset), explore whether the research could be conducted using synthetic data or the organization's own deployment data, and consider whether facilitating independent research serves the organization's governance interests even if it requires vendor negotiation

B. Refuse the researcher's request outright because the vendor agreement's confidentiality provisions create a legal obligation that cannot be overridden

C. Share the complete training dataset with the researcher without consulting the vendor because academic research is protected under the GDPR's research exemption

D. Share the training data only if the researcher agrees to keep the findings confidential and not publish them publicly

44. A governance professional is evaluating the organization's use of AI across its operations. The professional creates a matrix mapping each AI system against its risk classification, governance status, and the populations it affects. The matrix reveals that five different AI systems all affect the same population — hourly retail workers — but each system was independently assessed as low-to-moderate risk. The systems are: (1) AI shift scheduling, (2) AI performance monitoring, (3) AI task assignment, (4) AI break compliance tracking, and (5) AI attendance prediction. What governance insight does this portfolio-level analysis reveal?

A. The five systems collectively operate within normal governance parameters because each was independently assessed and classified appropriately based on its individual risk profile

B. The five systems should be consolidated into a single integrated system to reduce governance overhead and simplify the compliance landscape

C. The cumulative effect of five AI systems monitoring, scheduling, evaluating, and predicting the behavior of the same worker population creates a comprehensive surveillance environment whose aggregate impact on workers' autonomy, dignity, and well-being exceeds what any individual system assessment captured

D. The portfolio analysis is only relevant if all five systems were purchased from the same vendor, which would create vendor concentration risk

45. A governance professional is conducting a risk assessment for an AI system that will be deployed in an autonomous vehicle. The development team presents a risk matrix showing that all identified risks have been mitigated to "acceptable" residual risk levels. The professional notices that the risk matrix does not include any risks categorized as "catastrophic" in severity. For an autonomous vehicle system, what should the professional conclude?

A. The risk matrix is likely accurate because modern autonomous vehicle technology has advanced to the point where catastrophic failure modes have been engineered out of production systems

B. The absence of catastrophic severity risks indicates that the development team has effectively eliminated the most severe potential harms through robust engineering and safety design

C. The absence of catastrophic risks is not a concern because autonomous vehicle AI systems are classified as minimal risk under the EU AI Act and are not subject to risk assessment requirements

D. The absence of catastrophic severity risks in an autonomous vehicle context suggests the risk assessment may be incomplete — a system that controls a vehicle at high speeds in proximity to pedestrians has inherent catastrophic failure modes (collision, death, mass casualty) that must be identified, evaluated, and addressed even if they are low probability

46. A governance professional is reviewing the organization's AI system for customer creditworthiness assessment. The system was originally validated for personal loans up to €10,000. The business team has begun using it for business loans up to €500,000 without governance review. When questioned, the business team argues that "credit assessment is credit assessment regardless of the amount." What is the professional's MOST accurate response?

A. The professional should explain that the system was validated for a specific use case (personal loans up to €10,000) and applying it to business loans of 50x the value — with different risk profiles, different data characteristics, different regulatory requirements, and different consequences of error — constitutes a secondary use requiring independent validation and governance review

B. The business team is correct because the underlying credit assessment methodology is the same regardless of loan size or type

C. The professional's concern is valid only if the business loans are subject to different regulatory requirements than personal loans

D. The professional should approve the expanded use as long as the business team monitors the system's performance on business loans for the first 90 days

47. A governance professional is reviewing an AI system's testing results. The development team reports that the system achieved 94% accuracy on the test set. The professional asks how the test set was constructed. The development team explains that the test set was created by randomly sampling 20% of the available data, with the remaining 80% used for training. The professional identifies a potential issue. For which type of AI application is this standard random split MOST problematic?

A. An image classification system used for quality control in manufacturing, where product defect patterns are relatively stable over time

B. A time-series forecasting system used for financial market predictions, where random splitting can create temporal leakage by allowing the model to train on future data and be tested on past data

C. A text classification system used for email spam filtering, where spam patterns are generated from the same distribution as normal email

D. A recommendation system used for e-commerce product suggestions, where user preferences are relatively stable across the dataset

48. A governance professional receives a complaint from a customer who was denied credit by the organization's AI system. The customer requests an explanation. The governance professional works with the data science team to generate an explanation using a post-hoc interpretability tool (LIME/SHAP). The explanation identifies the top three contributing factors. However, the professional discovers that different interpretability tools produce different explanations for the same decision. What governance challenge does this inconsistency reveal?

A. The inconsistency indicates the AI model is defective and must be replaced with a model that produces consistent explanations across all interpretability tools

B. The inconsistency is expected and acceptable because different interpretability tools use different mathematical approaches and minor variations in explanation are normal

C. The inconsistency reveals a fundamental challenge in AI explainability — different methods can produce different and sometimes conflicting explanations for the same decision, raising questions about which explanation is "correct" and whether any post-hoc explanation method faithfully represents the model's actual reasoning process

D. The inconsistency can be resolved by selecting the interpretability tool that produces the most favorable explanation for the organization's compliance position

49. A governance professional is advising the board of directors on AI governance investment. The board asks: "How do we know our governance program is actually reducing risk, not just creating paperwork?" What type of evidence should the professional present to demonstrate governance effectiveness?

A. Outcome-oriented metrics such as declining trends in AI incidents and complaints, reduced time from incident detection to containment, decreasing percentage of deployed systems with identified governance gaps, and demonstrable cases where governance controls prevented foreseeable harms before they materialized

B. Activity metrics such as the number of impact assessments completed, the volume of governance documentation produced, and the total hours spent in governance committee meetings

C. Financial metrics such as the total governance budget and the number of full-time governance staff employed by the organization

D. Comparative metrics showing that the organization's governance program has more policies than its industry competitors

50. A governance professional is evaluating a proposal to deploy an AI system that predicts which patients in a hospital are at risk of developing pressure ulcers (bedsores). The system would alert nursing staff to high-risk patients so preventive measures can be taken. The system has been clinically validated with strong performance. However, the professional notes that the hospital's nursing staff is already severely understaffed. If the system generates alerts that nurses cannot respond to due to staffing constraints, what governance concern does this create?

A. The governance concern is limited to ensuring the system's alert interface is designed for maximum efficiency so nurses can process alerts quickly

B. The system should not be deployed in understaffed environments because AI systems require full staffing to function effectively

C. The staffing concern is an operational issue that falls outside the scope of AI governance and should be addressed by hospital administration

D. Deploying a system that generates alerts the organization cannot act upon creates a governance problem — the organization will have documented evidence of known patient risk that it failed to address, potentially increasing rather than decreasing liability, and the system's value depends on the organizational capacity to respond to its outputs

51. A governance professional is reviewing an AI system that generates automated responses to regulatory inquiries on behalf of the organization. The system drafts responses based on historical regulatory correspondence and the organization's compliance documentation. The professional identifies a critical governance risk that the compliance team has not considered. What is it?

A. The system may produce responses that are factually accurate but strategically disadvantageous — revealing more information than necessary, making inadvertent admissions, or framing the organization's position in ways that create legal exposure — and regulatory responses require human legal judgment that AI cannot provide

B. Automated responses to regulatory inquiries will save the organization significant legal costs by reducing the need for outside counsel

C. The system's responses may be flagged by regulators as AI-generated, which could result in automatic rejection of the organization's compliance submissions

D. The only risk is that the system may reference outdated regulations if its training data is not regularly updated with new regulatory developments

52. A governance professional is conducting a review of the organization's AI monitoring practices across all deployed systems. The review reveals that monitoring configurations vary widely — some systems have comprehensive monitoring with fairness metrics, drift detection, and performance tracking, while others have only basic uptime monitoring. The professional discovers that the monitoring depth does not correlate with the systems' risk classifications — some high-risk systems have minimal monitoring while some low-risk systems have comprehensive monitoring. What does this finding indicate?

A. The finding is acceptable because monitoring depth should be determined by the technical complexity of each system rather than by risk classification

B. The finding indicates that the monitoring team has appropriately prioritized the systems they find most technically interesting rather than following rigid governance hierarchies

C. The finding reveals a disconnect between the organization's risk classification framework and its monitoring implementation — monitoring depth should be proportionate to risk level, with high-risk systems receiving the most comprehensive monitoring, and the current misalignment means the organization's highest-risk systems may have the least governance visibility

D. The finding is only a concern if the monitoring configurations were set by the AI vendors rather than by the organization's own monitoring team

53. A governance professional is reviewing a proposal to use an AI system that analyzes employees' keystroke patterns, mouse movements, and screen activity to calculate "productivity scores." The system would run continuously during work hours on all company laptops. The professional must evaluate the proposal against the EU AI Act and GDPR. What is the MOST critical regulatory concern?

A. The system violates GDPR's data minimization principle because keystroke patterns and mouse movements constitute excessive data collection relative to the purpose of measuring productivity

B. The system requires registration in the EU database for high-risk AI systems before deployment because employee monitoring systems are classified as high-risk under Annex III

C. The system's processing of behavioral biometric data (keystroke dynamics) without explicit consent may violate GDPR's special category data provisions

D. The EU AI Act restricts AI-based emotion recognition in workplaces, and continuous behavioral monitoring that generates productivity scores based on typing and mouse patterns may fall within these restrictions — particularly if the system infers engagement, stress, or attention levels from behavioral biometric patterns

54. A governance professional is advising on the deployment of an AI system for predictive policing in a city with well-documented historical patterns of racially discriminatory enforcement. The police department argues that using recent data (last two years only) rather than historical data eliminates the bias concern because "recent enforcement is more equitable." The professional disagrees. Why?

A. Two years of data is insufficient for any enforcement patterns — but the deeper issue is that even recent enforcement data may reflect ongoing structural biases in policing patterns, resource allocation,

and community-police relationships that have not been fully remediated, making the recency of data an insufficient proxy for absence of bias

B. The police department's argument is correct because two years is the standard lookback period established by the NIST AI RMF for training data in law enforcement applications

C. The professional disagrees only because two years of data is too small a sample size for statistical validity in predictive policing applications

D. The professional's concern is limited to ensuring that the training data has been anonymized to remove personally identifiable information from the enforcement records

55. A governance professional is reviewing a deployed AI chatbot that provides legal information to the public about tenant rights. The chatbot was trained on legal texts from the organization's jurisdiction. A user in a different jurisdiction asks the chatbot about their tenant rights. The chatbot provides information based on its training jurisdiction's laws without disclosing that the information may not apply to the user's jurisdiction. What governance controls should have prevented this?

A. The chatbot should have been programmed to refuse to answer any questions from users outside the organization's jurisdiction to eliminate the risk of providing incorrect legal information

B. The chatbot should include jurisdiction detection and disclosure mechanisms that identify when a user may be in a different jurisdiction, clearly disclose that the information is specific to the training jurisdiction, and recommend that the user seek jurisdiction-specific legal resources

C. The governance concern is limited to adding a generic disclaimer that "laws vary by jurisdiction" to the chatbot's welcome message

D. The chatbot's training data should include legal texts from every jurisdiction worldwide to ensure comprehensive coverage of all possible user locations

56. A governance professional is evaluating an AI system that generates automated performance reviews for teachers based on student test scores, classroom observation data from video cameras, and parent satisfaction survey results. The system produces numerical performance ratings that directly influence compensation and continued employment. What is the MOST significant governance concern with this deployment?

A. The system's reliance on student test scores as a performance indicator may penalize teachers who work with disadvantaged student populations where factors beyond the teacher's control (poverty, food insecurity, unstable housing) significantly affect student performance

B. The system should not use parent satisfaction surveys because parents may provide biased feedback based on personal relationships rather than objective teaching quality

C. The system should only generate qualitative narrative reviews rather than numerical ratings because numbers create a false impression of measurement precision

D. The concern is limited to ensuring that the video observation cameras have sufficient resolution to accurately capture classroom interactions

57. A governance professional is asked to evaluate the governance implications of the organization deploying an AI system that monitors social media to identify potential employee candidates — proactively identifying individuals who demonstrate skills the organization needs, even though those individuals have not applied for a job. The system would compile candidate profiles from public social media, professional networking sites, and publicly available publications. What is the MOST fundamental governance question the professional should raise?

A. Whether proactively compiling profiles of individuals who have not expressed interest in employment with the organization — processing their personal data without their knowledge for a purpose they did not anticipate — complies with data protection principles including lawful basis, transparency, and purpose limitation

B. Whether the system can identify candidates with sufficient accuracy to justify the development and deployment costs

C. Whether the organization's HR department has the capacity to process the volume of candidate profiles the system would generate

D. Whether the social media platforms' terms of service permit automated data collection for recruitment purposes

58. A governance professional discovers that the organization's AI fraud detection system has been generating "risk scores" for individual customers and storing these scores in customer profiles visible to all customer-facing employees. Some employees have been using the risk scores to provide different levels of customer service — treating high-risk-scored customers with suspicion and providing lower-quality service. The fraud detection system's purpose is transaction monitoring, not customer service differentiation. What governance principle has been violated?

A. The principle of data minimization, because the risk scores should not be stored in customer profiles where they are accessible to employees who do not need them for fraud investigation

B. The principle of transparency, because customers have not been informed that their interactions are influenced by AI-generated fraud risk scores

C. The principle of safety, because the risk scores may be inaccurate and relying on them for customer service decisions creates a safety risk for the organization

D. Purpose limitation — the risk scores were generated for fraud detection, and their use to differentiate customer service quality constitutes secondary use beyond the system's intended purpose, creating both a governance violation and a fairness concern if the scores correlate with demographic characteristics

59. A governance professional is reviewing the organization's AI governance committee's composition and decision-making process. The committee includes representatives from legal, compliance, data science, product management, and risk management. However, the professional notes that the committee has never rejected or modified a proposed AI deployment — every proposal presented has been approved as submitted. What governance concern does this pattern raise?

A. The approval rate indicates the committee has established such effective pre-screening criteria that only well-governed proposals reach the committee

B. A 100% approval rate over time suggests the committee may not be functioning as an effective governance check — it may be rubber-stamping proposals rather than conducting genuine critical review, or the organizational culture may discourage committee members from raising objections that delay product launches

C. The approval rate is acceptable as long as the committee is documenting its deliberation process for each proposal

D. The governance committee should reject at least 20% of proposals to demonstrate it is functioning effectively as an oversight mechanism

60. A governance professional is advising an organization that has received a customer complaint about its AI system. The customer alleges that the AI system denied their insurance claim unfairly. The organization's investigation reveals that the AI system correctly applied the policy terms and the denial was accurate. However, the customer's experience of interacting with a fully automated system — receiving a denial with no human contact, no empathy, and no opportunity for real-time discussion — has caused significant distress. The professional argues this is still a governance concern even though the decision was correct. Why?

A. The professional is incorrect because the decision was accurate and governance is only concerned with the correctness of AI outputs, not the customer experience

B. The concern is limited to the organization's customer service department and has no connection to AI governance

C. The governance concern is that the customer experience of interacting with a fully automated denial process — without human contact, empathy, or real-time discussion — raises human-centricity and accountability concerns even when the decision is correct, because governance must consider how AI systems affect human dignity and autonomy, not just whether they produce technically accurate outputs

D. The professional's concern is only valid if the customer belongs to a protected demographic group under nondiscrimination law

61. An AI system used by a government welfare agency has been in operation for five years. During a routine governance review, the professional discovers that the agency has never conducted a post-deployment impact reassessment — the only impact assessment on file was conducted before the original deployment. During the five-year period, the agency's service population has changed significantly, new regulations have been enacted, and the system has been retrained three times. What governance obligation has been neglected?

A. Impact assessments are one-time pre-deployment documents and are not required to be updated during the system's operational life

B. Impact reassessments are only required when the system's risk classification changes from one tier to another

C. The agency is only required to update the impact assessment if a specific incident involving the system has been reported to the regulatory authority

D. Impact assessments must be reassessed when material changes occur — significant population changes, new regulatory requirements, and model retraining each constitute material changes that may alter the system's risk profile, and five years without reassessment means the current assessment does not reflect the system's actual operating context

62. A governance professional is reviewing the organization's approach to managing AI-related intellectual property. The organization uses several open-source AI models with different licenses. Some are under permissive licenses (Apache 2.0), some under copyleft licenses (GPL), and some under

restricted licenses with "responsible use" clauses. The organization has not tracked which models are used in which products. What governance risk does this create?

A. The organization may be unknowingly violating license terms — copyleft licenses may require the organization to open-source its own modifications, responsible use clauses may prohibit certain applications, and mixing incompatible licenses in a single product may create legal conflicts — all of which are invisible without systematic license tracking

B. Open-source license tracking is solely the responsibility of the organization's IT department and does not fall within the scope of AI governance

C. The risk is limited to the possibility that the open-source models may contain security vulnerabilities that proprietary models would not have

D. License tracking is only necessary for models that process personal data and is not required for models that process non-personal data

63. A governance professional is asked to evaluate a vendor's claim that its AI system is "ISO 42001 certified." The vendor presents a certificate issued by a consulting firm that helped the vendor implement the standard. The professional is skeptical. What is the basis for this skepticism?

A. ISO 42001 certification does not exist because the standard is still in draft form and has not been officially published

B. ISO 42001 certification must be issued by an accredited certification body through an independent audit process — a certificate from the consulting firm that helped with implementation is not a valid certification because the certifier must be independent from the implementer to ensure objectivity

C. ISO 42001 certification is valid only for organizations with more than 1,000 employees and the vendor may be too small to qualify

D. The vendor's ISO 42001 certification is automatically valid if the vendor also holds ISO 27001 certification because the two standards share the same requirements

64. A governance professional is reviewing the organization's AI system for automated customer complaint routing. The system categorizes complaints by topic and severity and routes them to appropriate departments. The professional discovers that the system consistently routes complaints from customers who write in formal, professional language to senior customer service representatives, while

routing complaints written in informal language to junior representatives. What governance principle is at issue?

A. The principle of transparency, because customers are not informed that the formality of their writing style affects which representative handles their complaint

B. The principle of data minimization, because the system should not analyze the linguistic style of customer communications for routing purposes

C. The principle of fairness — routing quality (access to senior vs. junior representatives) should be determined by the nature and severity of the complaint, not by the customer's writing style, which correlates with education level, socioeconomic status, and language proficiency

D. The principle of accountability, because no individual has been assigned responsibility for monitoring the routing algorithm's behavior

65. A governance professional is conducting a lessons-learned review after an AI incident. The incident involved an AI system that produced harmful outputs for three months before detection. The root cause analysis identified that the monitoring system was operational but the alert thresholds were set too high — the system was degrading gradually and the alerts never triggered because the degradation at each monitoring interval was within the threshold. What monitoring design improvement should the professional recommend?

A. Implementing trend-based alerting in addition to threshold-based alerting — tracking not just whether metrics cross absolute thresholds at any single point but whether metrics are moving in a concerning direction over time, enabling detection of gradual degradation that threshold-only monitoring misses

B. Lowering all alert thresholds by 50% to ensure earlier detection of any performance changes

C. Replacing automated monitoring with manual daily reviews by the data science team to catch subtle performance changes that automated systems miss

D. Adding more monitoring metrics to increase the total number of dimensions tracked by the monitoring system

66. A governance professional is evaluating the organization's AI incident reporting culture. An anonymous employee survey reveals that 40% of employees who discovered potential AI issues did not report them. The most common reasons cited were: "I didn't think it was serious enough to report," "I

didn't know who to report to," and "I was afraid of being blamed for causing problems." What governance improvement addresses ALL THREE barriers simultaneously?

A. Implementing a mandatory monthly AI incident report that every employee must submit, even if they have nothing to report

B. Assigning a dedicated AI incident reporter in each department who is solely responsible for filing all AI-related reports on behalf of department members

C. Increasing the financial penalties for employees who fail to report AI incidents within 24 hours of discovery

D. Establishing clear, accessible reporting channels with guidance on what constitutes a reportable concern (addressing the "not serious enough" barrier), communicating reporting procedures widely (addressing the "didn't know who" barrier), and creating a non-punitive reporting culture that protects reporters from blame (addressing the fear barrier)

67. A governance professional is reviewing a proposal for an AI system that would analyze court sentencing data to identify disparities in sentences for similar crimes. The system would flag cases where sentences deviate significantly from statistical norms. Proponents argue the system promotes justice by identifying potential bias. Opponents argue the system could be misused to pressure judges into conformity. What governance analysis should guide the evaluation?

A. The proposal should be rejected because any AI system that evaluates judicial decisions violates the constitutional principle of judicial independence

B. The governance analysis should evaluate both the beneficial potential (identifying sentencing disparities) and the risk potential (chilling judicial discretion), assess what safeguards could preserve the benefits while mitigating the risks (such as limiting the system to aggregate pattern analysis rather than individual case flagging), and determine whether the system's design can be constrained to prevent misuse

C. The proposal should be approved without conditions because identifying sentencing disparities is an unqualified public good that requires no governance scrutiny

D. The governance analysis should focus exclusively on the system's technical accuracy in identifying statistical outliers, because if the system is accurate, its governance implications are automatically acceptable

68. A governance professional is advising a healthcare organization that wants to deploy an AI diagnostic system trained on data from tertiary care hospitals. The system will be deployed in primary care clinics. The professional raises a concern about the difference between these clinical settings. What is the MOST significant governance implication of this deployment context mismatch?

A. The cost of deploying the system in primary care clinics is lower than in tertiary hospitals, which makes the business case more favorable and reduces governance concerns

B. Primary care clinics have less IT infrastructure than tertiary hospitals, creating technical deployment challenges that require additional engineering resources

C. Tertiary hospitals see more complex, advanced-stage conditions while primary care sees earlier-stage presentations — the system trained on tertiary data may not accurately diagnose conditions as they present in primary care, where symptoms are subtler, test availability is limited, and patient demographics may differ significantly

D. The deployment context mismatch only matters if the primary care clinics are in a different country from the tertiary hospitals where the training data was collected

69. A governance professional discovers that an AI system's training data includes data purchased from a data broker. The data broker certified that all data was "lawfully collected." However, recent media reports reveal that the data broker obtained some of its data through deceptive practices — telling consumers their data would be used for "academic research" when it was actually being sold commercially. What is the organization's governance exposure?

A. The organization has no exposure because it relied in good faith on the data broker's certification of lawful collection

B. The organization's exposure is limited to requesting a refund from the data broker for the fraudulently certified data

C. The exposure is limited to the AI vendor that integrated the data broker's data into the training pipeline

D. The organization may face governance exposure because the training data's foundation is compromised — data obtained through deceptive practices may lack valid consent, violating the lawful basis for processing, and models trained on this data could be subject to enforcement action including potential algorithmic disgorgement

70. A governance professional is reviewing the organization's AI system for automated content moderation on a platform used by children ages 8-14. The system blocks content containing certain keywords related to self-harm. However, the professional discovers that the keyword-based approach is blocking children who are trying to access mental health resources and support content that uses the same keywords in a supportive, educational context. What governance improvement is MOST appropriate?

A. Implementing context-aware content moderation that distinguishes between content that promotes self-harm and content that provides mental health support — ensuring that safety mechanisms do not inadvertently block the very resources that vulnerable children need, while maintaining protection against genuinely harmful content

B. Removing all keyword-based content blocking because the false positive rate for support content outweighs the benefit of blocking harmful content

C. Adding the mental health resource URLs to a whitelist that bypasses the content moderation system entirely for all content from those sources

D. Replacing the keyword-based system with human moderators who review all content posted to or accessed on the platform in real time

71. A governance professional is evaluating the organization's compliance with GDPR Article 22 for its deployed AI systems. The professional identifies one system — an automated insurance underwriting system — that makes coverage and pricing decisions without meaningful human involvement. The compliance team argues that the system is exempt from Article 22 because the decisions are "necessary for entering into a contract" with the customer (Article 22(2)(a) exception). The professional questions this interpretation. Why?

A. The "necessary for entering into a contract" exception does not apply because insurance underwriting decisions precede the contract and are not part of the contractual performance itself

B. Even when the contractual necessity exception applies, the organization must still implement suitable safeguards including the right to obtain human intervention, the right to express a point of view, and the right to contest the decision — the exception permits the automated processing but does not eliminate the safeguard requirements

C. The contractual necessity exception applies only to employment contracts and does not extend to insurance contracts under GDPR

D. Article 22 does not apply to insurance underwriting because insurance is regulated under sector-specific financial services legislation rather than GDPR

72. A governance professional is conducting a cross-functional review of how different departments within the organization use AI-generated content. The review reveals the following: Marketing uses AI to generate social media posts (published without human review), Legal uses AI to draft contract clauses (reviewed by attorneys before use), Customer Service uses AI to generate email responses (sent automatically without review), and HR uses AI to draft job descriptions (reviewed by recruiters before posting). What governance concern should the professional prioritize?

A. The HR department's use of AI for job descriptions, because employment-related content is classified as high-risk under the EU AI Act and requires the most governance attention

B. The Legal department's use of AI for contract drafting, because legal documents have the highest potential financial impact if they contain errors

C. The departments publishing AI-generated content without human review (Marketing and Customer Service), because unreviewed AI outputs sent to external audiences create the highest immediate risk of harm — including false claims, inappropriate content, privacy violations, and brand damage — that reviewed outputs do not

D. All four uses require identical governance controls because the risk of AI-generated content errors is the same regardless of whether human review occurs before publication

73. A governance professional is advising the board on the organization's AI governance program maturity. The board asks what distinguishes a "managed" governance program from an "optimizing" one. What is the KEY difference?

A. A managed program has standardized governance processes that are consistently applied and measured across the organization, while an optimizing program systematically learns from governance experiences — sharing insights across projects, adapting practices based on accumulated evidence, and proactively improving governance before problems arise rather than reacting to them

B. A managed program has 10-20 governance policies while an optimizing program has more than 50 governance policies covering every possible governance scenario

C. A managed program employs 5-10 governance professionals while an optimizing program employs more than 20 governance professionals across all business units

D. A managed program conducts annual governance audits while an optimizing program conducts quarterly governance audits

74. A governance professional is reviewing the organization's approach to obtaining consent for AI processing. The organization's mobile app includes a consent prompt that appears during account creation: "By creating an account, you consent to the use of AI to personalize your experience, improve our services, and for other purposes as described in our privacy policy." The privacy policy is a 47-page document. The professional identifies governance concerns with this consent mechanism. What is the MOST fundamental concern?

A. The consent prompt should include a more detailed description of each specific AI processing activity rather than referencing the 47-page privacy policy

B. The consent prompt violates GDPR requirements because it is not displayed in all official EU languages simultaneously

C. The consent prompt should require users to initial each page of the privacy policy to demonstrate they have read and understood the full document

D. The consent is not specific (bundled with account creation), not informed (meaningful information about AI processing is buried in a 47-page document), and potentially not freely given (consent is a condition of service) — failing multiple GDPR requirements for valid consent as a lawful basis for processing

75. A governance professional is asked to evaluate whether the organization should implement "AI governance by design" — embedding governance requirements into the AI development toolkit and CI/CD pipeline so that certain governance checks are automated and enforced at the technical level. For example, automated bias testing that must pass before code can be merged, or documentation templates that must be completed before deployment scripts execute. What is the MOST significant benefit of this approach?

A. AI governance by design eliminates the need for a governance committee because all governance decisions are automated through the technical pipeline

B. Embedding governance into the development pipeline ensures that baseline governance requirements are consistently applied to every AI system without depending on individual teams to remember and follow governance procedures — reducing the risk of governance gaps caused by human error, time pressure, or inconsistent process adherence

C. The approach is primarily useful for generating compliance documentation automatically, which reduces the administrative burden on the governance team

D. The approach only benefits organizations that use specific software development methodologies (Agile/Scrum) and provides no advantage for organizations using other development approaches

76. A governance professional is reviewing the organization's AI ethics principles. The principles state: "We are committed to developing AI that is fair, transparent, safe, and accountable." The professional argues that these principles, while well-intentioned, are insufficient as governance tools. What is the basis for this argument?

A. The principles should also include "privacy" and "human-centricity" to be complete, but otherwise they are adequate governance tools

B. The principles are too specific and should be replaced with a single overarching principle of "responsible AI" that encompasses all governance concerns

C. Principles without operationalization are aspirational statements, not governance tools — each principle must be translated into specific, measurable requirements, assigned to responsible individuals, integrated into decision-making processes, and monitored for compliance to function as actual governance rather than organizational rhetoric

D. The principles are adequate as governance tools because they establish clear organizational values that employees can reference when making AI-related decisions

77. A governance professional is evaluating the organization's use of an AI system that predicts employee flight risk — identifying employees likely to leave the organization within the next six months. The system uses data including email volume changes, calendar patterns, badge access frequency, and peer interaction metrics. The governance committee previously approved the system for "talent retention planning." The professional discovers that several managers are using the system's predictions to deny training opportunities and project assignments to employees flagged as flight risks — reasoning that investing in employees who are about to leave is wasteful. What governance violation has occurred?

A. The managers' use of flight risk predictions to deny opportunities constitutes a secondary use beyond the approved purpose ("talent retention planning"), and the denial of training and project opportunities based on predicted departure creates a self-fulfilling prophecy that may also constitute adverse employment action based on AI profiling

B. The governance violation is limited to the managers' failure to document their use of the flight risk predictions in their personnel decision records

C. No governance violation has occurred because the system was approved for "talent retention planning" and denying opportunities to flight-risk employees is a reasonable retention strategy

D. The governance violation is limited to the data science team's failure to update the model card with information about how managers are using the system's predictions

78. A governance professional is advising an organization that has discovered its AI vendor has been acquired by a Chinese technology company. The organization processes EU citizens' personal data through the vendor's AI system. The vendor's servers are located in the EU. The professional raises a governance concern. What is the MOST significant concern?

A. Chinese technology companies are prohibited from operating AI systems in the EU under the EU AI Act's foreign ownership restrictions

B. The concern is limited to verifying that the acquiring company will honor the existing vendor agreement's pricing terms and service level commitments

C. The concern is limited to ensuring that the vendor's servers remain physically located within the EU after the acquisition

D. The acquisition may trigger cross-border data transfer concerns — even if servers remain in the EU, the change of ownership to a non-EU entity may affect the data governance framework, may create new government access risks under the acquiring company's national laws, and may require reassessment of the transfer mechanisms and safeguards under GDPR

79. A governance professional is reviewing the organization's approach to AI system retirement. The organization's current practice is to simply shut down AI systems when they are no longer needed, without any formal decommissioning process. The professional discovers that a recently retired AI system's monitoring logs, incident reports, and model documentation were deleted when the system was shut down. What governance requirement has been violated?

A. Decommissioning is an informal process and no specific governance requirements apply to how organizations retire AI systems

B. Governance documentation — including monitoring logs, incident reports, model documentation, impact assessments, and decision records — must be retained for the period specified by the organization's retention policy and applicable regulations, even after the AI system is decommissioned, because these records may be needed for regulatory audits, litigation, incident investigation, or organizational learning

C. The only documentation that must be retained after decommissioning is the original impact assessment, and all other documentation can be deleted

D. Documentation retention requirements apply only to AI systems that were classified as high-risk under the EU AI Act

80. A governance professional is asked to prepare a summary for the AIGP certification exam. The professional must identify the ONE governance concept that connects every domain and every chapter of the AIGP Body of Knowledge. What concept is it?

A. The EU AI Act's risk classification system, which provides the legal framework that governs all AI activities worldwide

B. The NIST AI RMF's four core functions (Govern, Map, Measure, Manage), which provide the universal methodology for all AI governance activities

C. The principle that AI governance is a continuous, lifecycle-spanning, cross-functional responsibility that requires understanding what AI systems do, knowing the laws and standards that constrain them, governing their development with rigor, and maintaining vigilant oversight of their deployment and use — because every governance activity, from foundational principles to applied deployment governance, serves the goal of ensuring AI systems operate responsibly throughout their existence

D. ISO/IEC 42001's Plan-Do-Check-Act cycle, which provides the management system framework that all AI governance programs must follow

81. A governance professional discovers that the organization's AI system for customer service automation has been generating responses that inadvertently disclose other customers' personal information. The system occasionally includes fragments of previous customer conversations in its responses to new customers. Investigation reveals this is caused by the model memorizing specific interaction patterns during training. What is the MOST comprehensive governance response?

A. Add a disclaimer to all chatbot responses noting that the system may occasionally reference information from other interactions

B. Report the issue to the AI vendor and wait for a software patch before taking any additional action

C. Retrain the model with stronger regularization to reduce memorization and monitor for recurrence during the next quarterly review

D. Immediately implement output filtering to detect and block responses containing cross-customer data leakage, notify affected customers whose data was exposed, assess the scope of the privacy breach, evaluate regulatory notification obligations, retrain the model with privacy-preserving techniques, and implement ongoing monitoring specifically for memorization-related data leakage

82. A governance professional is reviewing the organization's AI governance committee meeting records. The records show that the committee meets monthly and discusses active AI projects. However, the professional notes that the meeting minutes do not document the rationale for governance decisions — they record only "approved" or "deferred" without explaining why. What governance principle does this documentation gap undermine?

A. The accountability principle — without documented rationale, the organization cannot demonstrate why governance decisions were made, cannot learn from past decisions, cannot defend its choices to regulators or in litigation, and cannot ensure consistency in future decisions on similar matters

B. The transparency principle only, because the meeting minutes should be published on the organization's public website for stakeholder review

C. No governance principle is undermined because the committee's approval or deferral decision is the only governance-relevant output of committee meetings

D. The documentation gap only matters if the committee has rejected a proposal, because rejections require justification while approvals do not

83. A governance professional is evaluating the organization's AI vendor ecosystem and discovers that the organization uses AI services from 12 different vendors, each with different data processing agreements, different security standards, and different governance commitments. No single individual or team has visibility across all 12 vendor relationships. What portfolio-level governance risk does this fragmentation create?

A. The fragmentation creates no governance risk because each vendor relationship is managed independently by the department that uses the service

B. The fragmentation is only a concern if two or more vendors are direct competitors, which would create antitrust issues

C. The fragmented vendor management prevents the organization from identifying correlated risks (multiple vendors with similar security vulnerabilities), inconsistent data governance (different processing standards for similar data across vendors), contractual gaps (missing provisions in some agreements that exist in others), and cumulative compliance exposure (the aggregate regulatory risk across all 12 relationships)

D. The fragmentation can be resolved by consolidating all AI services to a single vendor, which eliminates vendor management complexity

84. A governance professional is reviewing the organization's AI system for automated résumé screening. The system was originally deployed for entry-level technology positions. The HR department has gradually expanded its use to senior leadership positions, positions requiring security clearances, positions in regulated industries (healthcare, finance), and positions in different countries — all without governance review. Each expansion changes the system's deployment context in significant ways. What governance mechanism should have prevented this gradual scope expansion?

A. Technical access controls that limit the system to processing applications only for the originally approved job categories, combined with change management procedures requiring governance review before any expansion of the system's scope

B. An annual governance review that evaluates the system's current use cases against its original approved scope

C. A contractual limitation in the AI vendor agreement that restricts the system's use to the originally specified job categories

D. Quarterly user training reminders that inform HR staff about the system's approved use cases

85. A governance professional is conducting a comprehensive review of all AI-related risks across the organization. The review identifies a category of risk that is absent from the organization's risk register: the risk that the organization becomes excessively dependent on AI systems for critical business functions without adequate contingency planning. What is this type of risk called, and why is it significant?

A. AI dependency risk is significant because organizations that become critically dependent on AI systems without fallback procedures, manual processing capabilities, and business continuity plans may face operational paralysis if AI systems fail, vendors discontinue products, or regulatory changes require

system suspension — a risk that grows invisibly as AI becomes embedded in more business-critical processes

B. AI dependency risk is not a real governance concern because modern AI systems are sufficiently reliable that contingency planning is unnecessary

C. AI dependency risk is relevant only for organizations that use a single AI vendor for all their AI needs

D. AI dependency risk is addressed by the organization's existing IT disaster recovery plan and does not require separate consideration in the AI governance framework

86. A governance professional is reviewing the organization's data annotation process for AI training data. The organization uses a team of 20 annotators to label data for a healthcare AI system. The professional discovers that annotators receive no training on the clinical concepts they are labeling, have no access to clinical experts for clarification, and are evaluated primarily on annotation speed rather than accuracy. What governance risk does this annotation process create?

A. The risk is limited to annotation efficiency, which affects the project timeline but has no governance implications for the AI system's quality or fairness

B. The annotation process does not create governance risk because annotators' labels are verified by the AI model during training, which automatically corrects any labeling errors

C. The risk is limited to employee satisfaction among annotators, which is an HR concern rather than a governance concern

D. Untrained annotators labeling clinical data without expert guidance are likely to produce labels that contain systematic errors, misinterpretations, and inconsistencies — creating label bias that the AI model will learn as ground truth, potentially producing a healthcare system that reflects annotators' misunderstandings rather than clinical reality

87. A governance professional is reviewing the results of a bias audit conducted on the organization's AI lending system. The audit reveals that the system achieves statistical parity across racial groups for loan approval rates. However, the professional notices that approved minority applicants receive higher interest rates on average than approved non-minority applicants with similar credit profiles. The lending team argues that approval rate parity demonstrates fairness. How should the professional respond?

A. The lending team is correct because approval rate parity (demographic parity) is the definitive measure of fairness in lending and no additional analysis is needed

B. The professional should agree with the lending team's assessment but recommend periodic monitoring of interest rates as a secondary quality metric

C. Approval rate parity measures only one dimension of fairness — the interest rate disparity reveals a different dimension of unfair treatment that approval rate analysis alone cannot detect, because the system may approve minority applicants at equitable rates while systematically charging them more, creating discriminatory outcomes that require investigation and remediation

D. The interest rate disparity is an acceptable market-driven outcome because interest rates are determined by risk factors beyond the AI system's control

88. A governance professional is advising the organization on a proposed AI system that uses satellite imagery and public records to estimate residential property values. The system would be used by mortgage lenders for collateral valuation. The professional identifies that historical property valuation data — which the system would be trained on — reflects decades of racial discrimination in real estate (redlining, discriminatory appraisals, exclusionary zoning). What governance action is MOST essential before deployment?

A. The system should use only satellite imagery data (which does not directly encode racial discrimination) and exclude all historical valuation data from the training set

B. Evaluate the extent to which the training data's historical discrimination patterns are reproduced in the model's valuations, test whether the system systematically undervalues properties in historically redlined neighborhoods, implement fairness testing across geographic and demographic dimensions, and consider whether additional data sources or model adjustments can mitigate the perpetuation of historical discrimination

C. The historical discrimination in the training data is not a governance concern because the system is merely learning accurate market values as they currently exist

D. The system should be deployed with a disclosure to lenders that property valuations may be influenced by historical patterns, which transfers liability for any discriminatory impact to the lenders

89. A governance professional is evaluating the organization's AI governance program at the end of its first year. The program has successfully established policies, assigned roles, completed training, and begun monitoring deployed systems. The professional must recommend priorities for Year 2. Which priority would MOST advance the program's maturity?

A. Publishing more governance policies to address edge cases and unusual scenarios that the Year 1 policies did not cover

B. Shifting focus from establishing governance infrastructure to measuring governance effectiveness — implementing outcome metrics, conducting the first comprehensive audit, analyzing whether governance activities are actually reducing AI risk, and using findings to improve governance practices through an evidence-based continuous improvement cycle

C. Hiring additional governance staff to increase the team's capacity for handling the growing volume of governance reviews

D. Expanding the governance committee's membership to include representatives from every department in the organization

90. A governance professional is asked to provide a single-sentence summary of the MOST important lesson from the AIGP Body of Knowledge for organizations beginning their AI governance journey. What statement MOST accurately captures this lesson?

A. Know your AI systems — what they do, who they affect, what risks they create, and what laws apply — because every governance decision, from risk classification to monitoring design to incident response, flows from a clear understanding of the systems you are governing

B. Hire the most expensive governance consultants available because external expertise is more valuable than internal governance capability

C. Focus exclusively on EU AI Act compliance because it is the only regulatory framework that matters for AI governance globally

D. Implement the most restrictive governance controls possible for every AI system because over-governance is always preferable to under-governance

91. A governance professional is reviewing an AI system that recommends which students should receive academic intervention (tutoring, mentoring, counseling). The system was trained on data from schools where intervention was historically provided based on teacher referrals. Analysis reveals that teachers in the training data disproportionately referred students from affluent families for academic support, while students from lower-income families who needed support equally were less likely to be referred. The AI system has learned these patterns. What governance concept does this illustrate?

A. A technical performance issue that can be resolved by improving the AI model's accuracy in predicting which students need academic intervention

B. A data security concern because the training data contains sensitive information about students' academic performance and family income levels

C. Historical bias encoded in training data — the AI system perpetuates the inequitable referral patterns of human teachers rather than identifying students based on actual academic need, potentially widening educational achievement gaps rather than closing them

D. A privacy concern because the system processes student data that may require parental consent under education privacy laws

92. A governance professional is evaluating the organization's AI system for automated insurance claim adjudication. The system was deployed 18 months ago and has processed 500,000 claims. The professional requests the system's monitoring records and discovers that monitoring has been operational for only 6 of the 18 months — it was activated three months after deployment and was offline for two separate periods totaling nine months due to infrastructure issues. What is the governance significance of these monitoring gaps?

A. The monitoring gaps are a minor infrastructure concern that should be addressed by the IT team without governance involvement

B. The monitoring gaps create a governance compliance issue that can be resolved by generating estimated monitoring data for the unmonitored periods based on the available data

C. The monitoring gaps are only significant if a specific incident occurred during an unmonitored period

D. For 12 of 18 operational months, the organization has no governance visibility into the system's performance, fairness, drift, or security — creating compliance violations, accountability gaps for 333,000+ unmonitored decisions, and the inability to demonstrate that the system operated within acceptable parameters during the majority of its operational life

93. A governance professional is preparing the organization for its first AI governance audit. The professional reviews all documentation and discovers that while impact assessments, model cards, and test results are complete, there is no documentation showing how governance decisions were actually implemented — no evidence linking the impact assessment's recommended mitigations to actual controls deployed in the system, no verification that test findings led to remediations, and no confirmation that governance committee conditions were fulfilled before deployment. What does this gap reveal?

A. The organization has governance documentation (assessments, cards, tests) but lacks governance traceability — the ability to demonstrate that governance findings were acted upon, that conditions were fulfilled, and that identified risks were actually mitigated rather than merely documented

B. The gap is immaterial because the existence of impact assessments, model cards, and test results is sufficient evidence of governance for audit purposes

C. The gap can be resolved by having the governance committee retroactively sign off on all deployed systems, which creates the required implementation evidence

D. The gap is only relevant if the auditor specifically requests implementation evidence, and the organization should wait for the auditor's request before addressing it

94. A governance professional is evaluating a proposal to deploy an AI system that would analyze emergency department patients' electronic health records and vital signs to predict which patients are likely to experience cardiac arrest within the next four hours. The system would alert the clinical team to initiate preventive measures. The system has been clinically validated with strong predictive performance. However, the professional notes that the system's training data came from a hospital with a specific patient demographic and clinical protocol set. The deploying hospital serves a different patient population and uses different clinical protocols. What is the SINGLE most important governance question?

A. Whether the deploying hospital has the IT infrastructure to support the system's computational requirements

B. Whether the system has been validated on the deploying hospital's specific patient population and clinical context — because a system validated on one population using one set of clinical protocols may not perform equivalently for a different population with different disease prevalence, comorbidity patterns, and treatment approaches

C. Whether the system's vendor offers a service level agreement with guaranteed uptime above 99.9% for clinical environments

D. Whether the deploying hospital's nursing staff has been trained on how to respond to the system's cardiac arrest risk alerts

95. A governance professional is reviewing the organization's approach to managing bias in AI systems. The organization's current policy states: "All AI systems must be free from bias." The professional recommends revising this policy. What is the governance basis for this recommendation?

A. The policy is adequate because "free from bias" is a clear and achievable standard that all AI teams should be held to without exception

B. The policy should be revised to specify which types of bias are prohibited, because some forms of bias (such as favoring higher-quality products in a recommendation system) are acceptable and even desirable

C. "Free from bias" is an unachievable standard that may create perverse incentives — teams may avoid bias testing to maintain the appearance of compliance, or may game metrics to achieve nominal "bias-free" status, rather than engaging in honest assessment, transparent reporting, and continuous mitigation of identified biases

D. The policy is only problematic because it does not specify the mathematical definition of "bias" that should be used for measurement

96. A governance professional is advising the organization on a dispute between the data science team and the legal team. The data science team wants to include demographic data (race, gender, age) in the training data for a hiring AI system to enable the model to learn and correct for historical biases. The legal team argues that including protected characteristics in training data creates direct discrimination liability. Who is correct, and how should the professional resolve the dispute?

A. The legal team is absolutely correct — including any protected characteristic in AI training data constitutes direct discrimination that is illegal under all circumstances

B. The data science team is absolutely correct — including demographic data is the only way to detect and mitigate bias, and legal concerns are secondary to fairness objectives

C. Both raise valid points, and the governance professional should defer the decision to an external legal opinion without providing any governance guidance

D. Both teams raise legitimate concerns that must be balanced — including demographic data enables bias detection and mitigation (the data science team's valid point), but using protected characteristics as direct input features in decision-making creates discrimination risk (the legal team's valid point), and the resolution may involve using demographic data for bias testing and monitoring purposes while excluding it from the model's decision-making features

97. A governance professional is conducting a year-end review of the organization's AI governance program. The review reveals that governance has been consistently applied to new AI projects initiated during the year but that 12 AI systems deployed before the governance program existed ("legacy

systems") remain entirely ungoverned. These legacy systems include several that make consequential decisions affecting customers and employees. What should the professional recommend?

A. Legacy systems should be exempt from governance requirements because they were deployed before the governance program was established and retroactive governance is impractical

B. The professional should recommend a structured program to bring legacy systems into governance compliance — starting with a risk assessment of all 12 systems, prioritizing them by risk level, and developing a phased plan to bring each system into compliance with current governance standards, with highest-risk systems addressed first

C. Legacy systems should be immediately shut down and replaced with newly developed systems that comply with the governance framework

D. The professional should document the legacy systems' ungoverned status and include this information in the annual governance report without recommending any specific action

98. A governance professional is reviewing the organization's AI governance metrics dashboard. The dashboard tracks: number of impact assessments completed, number of model cards published, percentage of employees trained, number of governance committee meetings held, and number of policies published. The professional argues that the dashboard is measuring the wrong things. What type of metrics should replace or supplement these activity metrics?

A. Financial metrics measuring the cost savings generated by the AI governance program through prevented incidents and avoided regulatory fines

B. Outcome metrics measuring whether governance activities are actually effective — such as percentage of systems meeting fairness thresholds, incident detection time trends, proportion of identified risks successfully mitigated, and rate of governance findings that resulted in actual system improvements

C. Benchmarking metrics comparing the organization's governance activity levels against industry competitors

D. Staffing metrics measuring the ratio of governance professionals to deployed AI systems across the organization

99. A governance professional is preparing a presentation for new hires about the organization's AI governance program. The presentation must convey, in a single key message, why AI governance matters and why every employee plays a role. What message MOST effectively communicates both the importance and the universality of AI governance?

A. AI governance matters because regulatory penalties for non-compliance can reach €35 million or 7% of global turnover, and employees who fail to follow governance procedures may face personal liability

B. AI governance is exclusively the responsibility of the governance team and new hires should direct all AI-related questions to the dedicated governance email inbox

C. AI governance exists to protect the organization's competitive advantage by ensuring that proprietary AI models are not exposed to competitors through security breaches

D. AI governance ensures that the AI systems we build and use treat people fairly, operate safely, and remain accountable — and because AI touches every function in our organization, every employee has a role in identifying risks, following governance procedures, and speaking up when something doesn't seem right

100. Having completed Practice Exam 6, a candidate reflects on their preparation across six full-length examinations totaling 600 questions. What study approach is MOST likely to improve performance on the remaining four practice exams?

A. Memorize the answer keys from Exams 1-6 because the remaining exams will reuse a significant percentage of identical questions

B. Focus exclusively on the questions answered incorrectly without reviewing the explanations for correctly answered questions

C. Analyze error patterns across all six exams to identify whether weak areas reflect knowledge gaps (concepts not understood), application errors (concepts understood but applied incorrectly), or reading errors (questions misread or qualifiers missed) — and target study to the specific type of error rather than restudying all material uniformly

D. Abandon the practice exams and focus exclusively on rereading Part One chapters because reading is more effective than testing for exam preparation

Practice Exam 6: Answer Key and Explanations

1. C — Governance cannot be designed or prioritized without first understanding what needs to be governed. An inventory documents what each system does, who it affects, what data it processes, and a preliminary risk estimate — providing the foundation for every subsequent governance activity from risk classification to policy development to monitoring design.

2. A — Aggregate demographic statistics about a validation dataset (percentage breakdowns by gender, age, ethnicity) do not constitute personal data and can be shared without violating data protection principles. The vendor's response is likely an evasion, and the deployer needs this information to fulfill its independent governance obligations for a high-risk system.

3. D — The system's purpose has fundamentally changed from spam filtering (a technical utility function) to employee sentiment surveillance (monitoring employees' attitudes toward the company). This constitutes a secondary use triggering new privacy considerations, potential employment law violations, and possibly EU AI Act workplace monitoring provisions — all requiring comprehensive governance review.

4. B — Proportionality analysis is the governance framework for resolving principle conflicts. The committee must evaluate whether the patient safety benefit justifies the privacy intrusion, whether less invasive alternatives exist (direct patient engagement, pharmacy records), whether patients can meaningfully consent, and whether prediction accuracy warrants clinical action.

5. A — A system processing a population not represented in its training data may produce unreliable results for that population even while aggregate metrics remain stable. Gig economy workers have different income patterns, employment histories, and credit profiles than the training population. Disaggregated analysis for this specific segment is essential to detect hidden performance or fairness issues.

6. C — Governance can be appropriately accelerated without being abandoned. A rapid risk assessment identifies minimum essential controls (performance validation, basic fairness check, human oversight, monitoring, rollback capability) that can be implemented within the compressed timeframe, with the full governance review scheduled for completion within 30 days.

7. D — Without a centralized AI procurement process and inventory, governance cannot ensure consistent vendor assessment, contractual protections, data governance standards, or compliance

requirements across the organization. The three independent purchases with inconsistent terms demonstrate a structural governance gap that enables ungoverned AI acquisition.

8. B — Biometric access control verifies identity by matching a fingerprint against enrolled records — this IS identification through biometric data processing. Fingerprint data constitutes special category data under GDPR Article 9 regardless of whether the stated purpose is "access control" or "identification." The DPIA's conclusion fundamentally mischaracterizes the nature of the processing.

9. C — A change of vendor ownership could expose the organization's data to a competitor, alter the vendor's data processing practices, or transfer data to an entity with different privacy standards and legal obligations. Without contractual protections addressing change of ownership, the organization has no remedy if these risks materialize.

10. A — Real-time processing of facial images to estimate age constitutes personal data processing under GDPR even without storage. The processing involves analyzing a natural person's biometric features to infer a personal characteristic (age). GDPR's definition of processing includes any operation performed on personal data, not only storage.

11. D — The data science team leader has an inherent conflict of interest as the person responsible for the systems that may have caused the incident. Incident severity classification and response decisions should involve cross-functional authority — including governance, legal, compliance, and business stakeholders — to ensure objectivity and comprehensive response.

12. B — Data drift is a leading indicator that often precedes performance and fairness degradation. The appropriate response is proactive: investigate the cause, increase monitoring frequency, and prepare retraining plans. Waiting for performance impact means waiting for harm to materialize before acting, defeating the purpose of drift monitoring.

13. A — The vendor is using employee contact data collected for a specific purpose (system authentication) for a different purpose (marketing) without authorization. This likely violates data protection principles regarding purpose limitation and constitutes unauthorized processing of employee personal data that the organization — as the controller — should address.

14. C — Single-axis fairness testing (gender only) can miss disparities that emerge at the intersection of protected characteristics. A system may appear fair for women overall and for racial minorities overall

while producing significantly worse outcomes for women of a specific racial group. Intersectional testing is essential for comprehensive fairness evaluation.

15. D — Before evaluating technical accuracy or bias, the committee must determine whether the constructs being measured can be validly assessed from voice patterns at all. If "confidence," "trustworthiness," and "leadership potential" cannot be scientifically measured through vocal analysis, the system causes harm regardless of its technical precision — it makes decisions based on invalid measurements.

16. B — A technical fix addresses the immediate classification error but does not address why monitoring failed to detect the mass removal, whether affected users should be notified and their content restored, whether the system has systemic gaps in distinguishing documentation of violence from violent content, or what governance improvements would prevent recurrence.

17. A — Risk-proportionate oversight calibrates the oversight mechanism to the consequences of error. For lower-acuity classifications where delays are less harmful, immediate AI action with statistical review (Approach B) is appropriate. For high-acuity classifications where errors are life-threatening, human review before action (Approach A) is essential despite the time cost.

18. C — Indefinite retention of training data containing personal data may violate GDPR's data minimization and storage limitation principles. Personal data should be retained only for as long as necessary for a specified purpose. "In case it might be needed" is not a sufficiently defined retention purpose to justify indefinite storage.

19. D — "AI research and development" is a specific, limited purpose that almost certainly does not encompass commercial audiobook production and distribution. Commercial deployment constitutes a fundamentally different purpose that exceeds the scope of the original licensing authorization, creating both purpose limitation and intellectual property governance concerns.

20. A — Customer segmentation can produce discriminatory outcomes even without individual "decisions." If the system systematically places customers from certain demographic groups into less favorable segments, those customers may receive inferior service, fewer offers, or different pricing — constituting discriminatory treatment that fairness monitoring should detect.

21. B — Post-deployment activities (monitoring, maintenance, incident response, reassessment) continue for the entire operational life of the system — potentially years or decades. These sustained

costs typically exceed the one-time pre-deployment investment. An 80/20 split favoring pre-deployment underinvests in the longest and most consequential governance phase.

22. D — Public availability does not resolve purpose limitation. Individuals posted on social media for social communication, not AI training. Using that content for a fundamentally different purpose may not be compatible with the original purpose, regardless of the data's visibility. This is the most fundamental concern because it questions the lawful basis for the entire training dataset.

23. C — Protected characteristics and cultural contexts differ between jurisdictions. Bias testing valid in France (evaluating gender, ethnicity, religion under French nondiscrimination law) may not address the relevant fairness dimensions in Saudi Arabia, and vice versa. Each deployment requires jurisdiction-specific calibration.

24. A — The system's intended purpose was defined too narrowly, omitting arithmetic verification — a fundamental requirement for any invoice processing system. This design-phase governance gap allowed the system to automate approval without performing basic quality checks that would catch mathematical errors, resulting in €400,000 in preventable overpayments.

25. B — Without previous model versions, the organization cannot determine whether newly discovered issues existed before or were introduced by an update, cannot roll back to a known-good version if an update causes problems, and cannot investigate incidents requiring comparison of system behavior across versions. Version management is essential for governance continuity.

26. D — A standard risk classification framework may not adequately account for the catastrophic cascading consequences specific to nuclear environments. A missed failure prediction in a nuclear power plant carries consequences — radiation exposure, environmental contamination, public safety — of an entirely different magnitude than in standard industrial settings, warranting elevated risk classification.

27. A — Open-sourcing permanently eliminates trade secret protection and removes the organization's ability to control downstream uses. Once released, the model can be used for surveillance, discrimination, or other harmful applications without recourse. The organization must evaluate whether the transparency benefits outweigh these irreversible consequences before release.

28. C — 95% completion with low comprehension and behavioral change indicates the training measures participation, not learning. The program needs redesign focused on practical competence —

ensuring employees can identify AI systems they interact with, understand governance policies relevant to their role, and know how to report concerns.

29. D — Automated employment references that include information affecting future employment prospects constitute solely automated decision-making with significant effects under GDPR Article 22. The system processes personal data to generate an output that directly influences the individual's ability to secure future employment without meaningful human involvement.

30. B — The logging gap means the organization cannot demonstrate the basis for 200,000 decisions, cannot verify fairness or accuracy, cannot respond to customer complaints or regulatory inquiries about specific determinations, and may have frozen customer accounts based on completely unauditable AI outputs — creating severe compliance, accountability, and liability exposure.

31. A — Synthetic data generated from real patient records requires governance oversight. The generation process may not fully anonymize, the synthetic samples may closely resemble specific patients, and the fidelity must be validated to ensure the downstream model learns genuine patterns rather than artifacts. "Not real data" does not mean "no governance needed."

32. C — Historical transcripts containing expired promotional offers are presented as current information by the chatbot. The system lacks mechanisms to distinguish current from outdated offers, creating false representations to customers that may constitute deceptive practices. This is a training data governance issue requiring temporal data management.

33. B — Patent offices in most jurisdictions require a natural person as the inventor. The organization should identify the human researchers who designed the AI system, defined its research parameters, or interpreted its results as the inventors. Documentation of the human creative contribution is essential to support the patent application.

34. D — The three-week notification delay demonstrates the consequences of inadequate vendor agreements and the absence of ongoing vendor monitoring. Vendor agreements must include specific breach notification timeframes, and organizations must implement continuous vendor monitoring to detect security, financial, and governance changes during the contract period.

35. A — When a system is producing ongoing harm to a protected group, the containment principle requires stopping the harm first. The default response is to suspend the system or implement manual

review while remediation is developed, rather than allowing continued discriminatory outcomes during a potentially extended fix development period.

36. C — The system creates a ceiling effect by directing development resources based on current position rather than individual potential. An employee in a junior role expressing leadership aspirations is systematically directed away from leadership training — reinforcing the existing hierarchy rather than supporting equitable career development based on individual goals and capability.

37. D — Regardless of marketing language ("workplace wellness tool"), analyzing employees' emotional states during work activities falls within the EU AI Act's workplace emotion recognition restrictions. The regulatory classification is determined by what the system does (emotion recognition in a workplace), not by how the vendor characterizes it.

38. B — A single reviewer processing all AI-generated scores will experience reviewer fatigue, may develop automation bias (uncritically accepting AI outputs over time), and creates a single point of failure. Effective oversight requires multiple reviewers, rotation schedules, independent spot-checking, and metrics tracking whether genuine evaluation is occurring.

39. A — Model cards must be updated when significant changes occur. Retraining alters the model's learned patterns, making the original model card inaccurate as a description of the current system. An outdated model card is inadequate for deployer communication, regulatory compliance, and audit purposes.

40. C — Without management review, senior leadership does not systematically evaluate whether the governance program achieves its objectives, whether resources are adequate, whether policies need updating, and whether the organization's AI risk posture is acceptable. Strategic governance decisions are not being made, leaving the program without executive direction.

41. B — The FIRST concern is lawful basis — whether the organization's privacy notice authorized the use of complaint data for AI training purposes. If the privacy notice only authorized processing for "complaint resolution and service improvement," using complaint data for AI model training may require a separate lawful basis or privacy notice update before any processing begins.

42. D — "No longer needed" is too vague to implement. The policy lacks specific triggers for initiating deactivation, detailed procedures for executing it (stakeholder notification, transition planning, data

disposition), and post-deactivation requirements (documentation archival, lessons learned). Without these specifics, the policy provides no actionable guidance.

43. A — The professional should explore whether the vendor agreement actually prohibits sharing aggregate statistics or anonymized subsets (rather than complete raw data), whether the research could use synthetic or deployment data instead, and whether facilitating independent research serves governance interests. Creative problem-solving within constraints is preferable to reflexive refusal.

44. C — Five AI systems collectively monitoring, scheduling, evaluating, tracking, and predicting the behavior of the same worker population create a comprehensive surveillance environment. The cumulative impact on workers' autonomy, dignity, and well-being exceeds what any individual system's risk assessment captured — demonstrating why portfolio-level analysis matters.

45. D — An autonomous vehicle operating at high speeds near pedestrians has inherent catastrophic failure modes (collision causing death or mass casualty) that must appear in any complete risk assessment. Their absence suggests the assessment is incomplete — either failing to identify these scenarios or inappropriately categorizing them below catastrophic severity.

46. A — The system was validated for personal loans up to €10,000. Business loans of €500,000 have different risk profiles, data characteristics, regulatory requirements, and error consequences. Applying a narrowly validated system to a materially different use case without independent validation and governance review constitutes unauthorized secondary use.

47. B — For time-series data, random splitting creates temporal leakage — the model may train on data from future time periods and be tested on past data, artificially inflating performance because the model has "seen the future." Time-series applications require chronological splitting to maintain the temporal integrity of the evaluation.

48. C — Different interpretability methods producing different explanations for the same decision reveals a fundamental challenge: post-hoc explanation methods may not faithfully represent the model's actual reasoning process. This raises questions about which explanation is "correct" and whether any generated explanation can be relied upon for regulatory compliance or individual rights fulfillment.

49. A — Outcome metrics demonstrate governance effectiveness: declining incident trends show risk reduction, faster detection-to-containment times show improved response capability, decreasing

governance gaps show expanding coverage, and documented harm prevention shows concrete value. Activity metrics measure effort; outcome metrics measure results.

50. D — Deploying a system that generates alerts the organization cannot act upon creates a paradoxical liability increase — the organization has documented evidence of known patient risk that it systematically failed to address. The system's value depends entirely on the organizational capacity to respond to its outputs. Without that capacity, the system creates risk rather than reducing it.

51. B — AI-generated responses to regulatory inquiries may be factually accurate but strategically disadvantageous — revealing more information than necessary, making inadvertent admissions, or framing positions in ways that create legal exposure. Regulatory correspondence requires human legal judgment about what to say, how to say it, and what to omit.

52. C — Monitoring depth should be proportionate to risk level. When high-risk systems have minimal monitoring while low-risk systems have comprehensive monitoring, the organization's highest-risk systems have the least governance visibility. This misalignment means the most consequential potential failures are the least likely to be detected.

53. D — The EU AI Act restricts emotion recognition in workplaces. Continuous behavioral monitoring generating productivity scores from keystroke dynamics and mouse patterns may fall within these restrictions — particularly if the system infers engagement, stress, or attention levels from behavioral biometric patterns, which constitutes emotion-adjacent workplace monitoring.

54. A — Even recent enforcement data may reflect ongoing structural biases that have not been fully remediated. Policing patterns, resource allocation, community-police relationships, and enforcement priorities evolve slowly. Two years of data recency does not guarantee the absence of structural bias — it merely guarantees the bias is recent.

55. B — The chatbot should detect when a user may be in a different jurisdiction, disclose that its information is jurisdiction-specific, and recommend jurisdiction-appropriate resources. Providing legal information from the wrong jurisdiction without disclosure creates a concrete risk that users will rely on inapplicable law to make important decisions.

56. C — Using student test scores as a teacher performance indicator penalizes teachers who work with disadvantaged populations where factors beyond the teacher's control (poverty, food insecurity, housing

instability) significantly affect academic outcomes. The AI system embeds this structural inequity into its performance ratings, potentially punishing teachers who serve the most vulnerable students.

57. A — Proactively compiling profiles of individuals who have not applied for employment raises fundamental data protection questions: what is the lawful basis for processing their data without their knowledge, how does this comply with transparency requirements, and is recruitment sourcing a purpose these individuals would reasonably expect when posting on social media or publishing professionally?

58. D — The risk scores were generated for fraud detection and their use to differentiate customer service quality constitutes secondary use beyond the intended purpose. This violates purpose limitation and creates a fairness concern if the scores correlate with demographic characteristics — meaning certain customer groups systematically receive inferior service based on an AI determination made for a different purpose.

59. B — A 100% approval rate over time suggests the committee may not be functioning as an effective governance check. It may be rubber-stamping proposals, or organizational culture may discourage members from raising objections that delay product launches. Genuine critical review should occasionally result in modifications, conditions, or deferrals.

60. C — Governance must consider how AI systems affect human dignity and autonomy, not just whether they produce technically accurate outputs. A fully automated denial process without human contact, empathy, or real-time discussion raises human-centricity concerns — the experience of receiving a consequential denial from an impersonal system matters even when the decision is correct.

61. D — Impact assessments must be reassessed when material changes occur. Significant population changes, new regulations, and model retraining each alter the system's risk profile. Five years without reassessment means the current assessment — based on conditions that no longer exist — does not reflect the system's actual operating context.

62. A — Untracked open-source licenses create invisible legal risk. Copyleft licenses may require open-sourcing organizational modifications, responsible use clauses may prohibit certain applications, and incompatible licenses in a single product may create legal conflicts. Without systematic tracking, these risks are completely invisible to governance.

63. B — ISO management system certification must be issued by an accredited certification body through an independent audit process. A certificate from the consulting firm that helped implement the standard lacks the independence required for valid certification — the certifier must be a separate entity from the implementer to ensure audit objectivity.

64. C — Routing quality based on writing style rather than complaint substance creates inequitable service. Writing formality correlates with education level, socioeconomic status, and language proficiency — meaning the system effectively provides better service to more educated and affluent customers while providing inferior service to those who communicate informally.

65. A — Trend-based alerting detects gradual degradation by tracking whether metrics are moving in a concerning direction over time, even when no single data point crosses an absolute threshold. This complements threshold-based alerting and catches the specific failure mode demonstrated in this incident — slow degradation that threshold-only monitoring misses.

66. D — Clear reporting guidance addresses the "not serious enough" barrier by defining what constitutes a reportable concern. Widely communicated procedures address the "didn't know who" barrier by making reporting channels accessible. A non-punitive culture addresses the fear barrier by protecting reporters from blame. All three barriers require simultaneous intervention.

67. B — The governance analysis should evaluate both the beneficial potential (identifying sentencing disparities) and the risk potential (chilling judicial discretion), assess safeguards that could preserve benefits while mitigating risks (such as aggregate pattern analysis rather than individual case flagging), and determine whether the design can prevent misuse.

68. C — Tertiary hospitals see complex, advanced-stage conditions while primary care sees earlier presentations with subtler symptoms. A system trained on tertiary data may not accurately diagnose conditions as they present in primary care, where symptoms are less pronounced, available diagnostics differ, and patient demographics may diverge significantly from the training population.

69. D — Data obtained through deceptive practices lacks valid consent, compromising the lawful basis for any processing built upon it. Models trained on unlawfully collected data could face enforcement action including potential algorithmic disgorgement. The data broker's certification does not shield the organization from the underlying consent deficiency.

70. A — Context-aware moderation distinguishes between content that promotes self-harm and content providing mental health support. Safety mechanisms must not block the very resources vulnerable children need. The system must protect against harmful content while ensuring access to supportive, educational resources that use similar terminology.

71. B — Even when the contractual necessity exception under Article 22(2)(a) applies, the organization must still implement suitable safeguards — including the right to obtain human intervention, express a point of view, and contest the decision. The exception permits automated processing but does not eliminate the safeguard requirements that protect the data subject.

72. C — Marketing and Customer Service publish AI-generated content without human review, creating the highest immediate risk of external harm — false claims, inappropriate content, privacy violations, and brand damage reaching customers directly. The departments with human review (Legal and HR) have a built-in quality control layer that catches errors before publication.

73. A — A managed program has standardized processes consistently applied and measured. An optimizing program goes further by systematically learning from governance experiences — sharing insights across projects, adapting based on evidence, and proactively improving before problems arise. The key difference is organizational learning and continuous improvement.

74. D — The consent fails multiple GDPR requirements: it is not specific (bundled with account creation rather than granular for each processing purpose), not informed (meaningful AI processing information is buried in a 47-page document that no reasonable person reads in full), and potentially not freely given (consent is effectively required to use the service).

75. B — Embedding governance into the development pipeline ensures baseline requirements are consistently applied without depending on individual teams to remember procedures. This reduces governance gaps caused by human error, time pressure, competing priorities, or inconsistent adherence — making governance a system property rather than a behavior dependent on individual compliance.

76. C — Principles without operationalization are aspirational statements, not governance tools. Each principle must be translated into specific, measurable requirements, assigned to responsible individuals, integrated into decision-making processes, and monitored for compliance. Without this operationalization, principles remain organizational rhetoric rather than functioning governance mechanisms.

77. A — The managers' use of flight risk predictions to deny opportunities constitutes secondary use beyond the approved purpose of "talent retention planning." Denying training and assignments creates a self-fulfilling prophecy (employees denied opportunities are more likely to leave) and may constitute adverse employment action based on AI profiling — which is the opposite of retention.

78. D — Even if servers remain in the EU, ownership by a non-EU entity may affect data governance through new government access obligations under the acquiring company's national laws, may alter the corporate governance and data access structure, and may require reassessment of GDPR transfer mechanisms and safeguards to ensure continued data protection adequacy.

79. B — Governance documentation must be retained for the period specified by retention policies and applicable regulations even after system decommissioning. These records may be needed for regulatory audits, litigation, incident investigation, or organizational learning — deleting them upon shutdown eliminates the organization's ability to account for the system's historical operation.

80. C — The connecting concept across every domain is that AI governance is continuous, lifecycle-spanning, and cross-functional — requiring understanding of AI systems, knowledge of laws and standards, rigorous development governance, and vigilant deployment oversight. Every governance activity serves the goal of ensuring AI systems operate responsibly throughout their existence.

81. D — The comprehensive response addresses all dimensions: immediate containment (output filtering), affected individual notification (privacy breach), scope assessment (how extensive is the exposure), regulatory evaluation (notification obligations), technical remediation (privacy-preserving retraining), and ongoing prevention (specific memorization monitoring). A single-dimension response leaves other harms unaddressed.

82. A — Without documented rationale, the organization cannot demonstrate why decisions were made, cannot learn from past decisions, cannot defend choices to regulators, and cannot ensure consistency. Accountability requires not just recording outcomes but recording the reasoning that produced them.

83. C — Fragmented vendor management prevents identification of correlated risks across vendors, creates inconsistent data governance standards, leaves contractual gaps in some agreements that exist in others, and obscures cumulative compliance exposure. Portfolio-level vendor visibility is essential for identifying these cross-vendor governance risks.

84. B — Technical access controls limiting the system to approved job categories prevent unauthorized scope expansion at the system level. Change management procedures requiring governance review before expansion prevent it at the process level. Together, these mechanisms ensure that each deployment context change receives appropriate governance evaluation.

85. A — AI dependency risk grows invisibly as AI becomes embedded in more critical functions. Organizations that lack fallback procedures, manual processing capabilities, and business continuity plans face operational paralysis if systems fail, vendors discontinue products, or regulatory changes require suspension. This risk requires explicit identification and mitigation in the governance framework.

86. D — Untrained annotators labeling clinical data without expert guidance produce labels containing systematic errors and misinterpretations. The AI model learns these flawed labels as ground truth, potentially creating a healthcare system that reflects annotators' misunderstandings rather than clinical reality — embedding label bias at the foundation of the system's knowledge.

87. C — Approval rate parity measures only one dimension of fairness. The interest rate disparity reveals discriminatory treatment that approval analysis cannot detect — the system may approve minority applicants equitably while systematically charging them more. Both the approval decision and the terms of approval must be evaluated for fairness.

88. B — The most essential action is evaluating whether historical discrimination is reproduced in the model's valuations, testing for systematic undervaluation in historically redlined neighborhoods, implementing fairness testing across geographic and demographic dimensions, and exploring mitigation strategies. Simply excluding historical data or deploying with a disclaimer is insufficient.

89. B — Year 1 established governance infrastructure. Year 2 should shift to measuring effectiveness — implementing outcome metrics, conducting the first comprehensive audit, analyzing whether activities actually reduce risk, and using findings to improve practices. This evidence-based continuous improvement cycle advances the program from "established" to "effective."

90. A — Every governance decision — risk classification, monitoring design, compliance assessment, incident response — flows from understanding what AI systems do, who they affect, what risks they create, and what laws apply. Without this foundational understanding, governance activities lack direction, prioritization, and relevance.

91. C — The AI system has learned teachers' inequitable referral patterns rather than identifying students based on actual academic need. This perpetuates historical bias: students from lower-income families who need support are systematically overlooked, potentially widening achievement gaps that the intervention system was designed to close.

92. D — For 12 of 18 months, the organization has no governance visibility into the system's operation. This creates compliance violations (no monitoring records for a majority of operational time), accountability gaps (333,000+ decisions cannot be audited), and the inability to demonstrate that the system operated within acceptable parameters during most of its life.

93. A — The organization has governance documentation but lacks governance traceability — the ability to demonstrate that findings led to actions, conditions were fulfilled, and risks were actually mitigated. Without this linkage, documentation proves only that assessments were performed, not that they influenced the system's actual governance.

94. B — A system validated on one patient population using specific clinical protocols may not perform equivalently for a different population with different disease prevalence, comorbidity patterns, and treatment approaches. For a cardiac arrest prediction system where errors are life-threatening, deployment-context validation is the single most critical governance question.

95. C — "Free from bias" is unachievable — multiple fairness definitions are mathematically incompatible, all data contains some bias, and an impossible standard incentivizes avoidance of testing or metric gaming. Effective policy requires documented assessment, transparent reporting, and continuous mitigation rather than an impossible absolute standard.

96. D — Both teams raise legitimate concerns. Including demographic data enables bias detection (data science's valid point), but using protected characteristics as decision features creates discrimination risk (legal's valid point). The resolution involves using demographic data for testing and monitoring while excluding it from decision-making features — a governance design that serves both objectives.

97. B — Legacy systems making consequential decisions without governance represent the organization's highest-risk AI portfolio segment. A structured program should assess all 12 systems, prioritize by risk, and develop a phased compliance plan — starting with highest-risk systems. Neither blanket exemption nor immediate shutdown is appropriate.

98. B — Activity metrics (assessments completed, policies published) measure governance effort. Outcome metrics (systems meeting fairness thresholds, incident trends, risk mitigation rates) measure governance effectiveness. The board's question — "is governance reducing risk?" — can only be answered by outcome metrics that demonstrate whether activities produce results.

99. D — The message connects AI governance to its purpose (treating people fairly, operating safely, maintaining accountability), establishes universality (AI touches every function), and empowers action (identifying risks, following procedures, speaking up). This framing is more motivating and actionable than citing penalties, claiming exclusivity, or focusing on competitive advantage.

100. C — Analyzing error patterns identifies whether weak areas stem from knowledge gaps (concepts not understood), application errors (concepts understood but misapplied), or reading errors (questions misread). Each error type requires a different study approach, making targeted remediation more effective than uniform restudying.