EU AI Act Compliance Checklist

A practical checklist to assess AI systems, identify compliance gaps, and strengthen your AI governance framework.

1. AI System Inventory

A complete AI system inventory is the foundation of EU AI Act compliance. Before an organization can classify risk or assign obligations, it must know which AI systems exist, where they are used, who owns them, and whether they are internally developed or supplied by a third party. The inventory should cover traditional machine learning models, generative AI tools, AI agents, embedded AI features, automated decision systems, and AI-enabled workflow applications.

1.1 Identify AI Systems

List all AI models, AI agents, and AI-powered applications used across the organization. This should include systems used directly by employees as well as AI capabilities embedded inside larger platforms, such as HR systems, finance tools, CRM applications, customer support platforms, productivity suites, security tools, and analytics dashboards.

- **Record the system name:** Capture the product, model, workflow, agent, or application name used by the business.
- **Identify the business function:** Note whether the system supports HR, Finance, Customer Support, Marketing, Operations, Legal, Compliance, IT, or another area.
- **Assign a system owner:** Name the accountable business owner, technical owner, and compliance contact where possible.

- **Document the purpose:** Explain what the AI system is intended to do, such as screening resumes, generating customer replies, detecting fraud, forecasting demand, or summarizing contracts.
- **Capture lifecycle status:** Mark whether the system is in pilot, production, retired, under procurement, or under redesign.

Example: An HR team may use an AI-enabled recruitment platform to rank applicants against job descriptions. The inventory entry should identify the platform name, the HR owner, the recruitment workflow it supports, the vendor, the input data used, and whether the tool influences hiring decisions or only provides administrative support.

1.2 Map AI Usage

After identifying AI systems, map where and how they are used. Usage mapping helps determine whether an AI system affects employees, customers, consumers, regulated decisions, safety-related processes, or fundamental rights. This step is especially important because EU AI Act classification often depends on the use case, not only on the underlying technology.

- **Map business areas:** Document AI usage across HR, Finance, Customer Support, Marketing, Operations, IT, Legal, Procurement, Risk, and Compliance.
- **Identify user groups:** Note whether the system is used by employees, managers, customers, contractors, vendors, or external users.

- **Record decision impact:** Specify whether the AI output is advisory, automated, semi-automated, or used as a direct decision input.
- **Identify development source:** Mark whether the AI system is internally developed, customized from an open-source model, purchased from a third-party vendor, or embedded in a commercial platform.
- **Document data inputs and outputs:** Capture the type of data processed and the nature of the AI output, such as scores, recommendations, generated text, risk alerts, predictions, classifications, or automated actions.

Example: A customer support chatbot may be considered limited risk if it interacts with users and provides automated responses, creating transparency obligations. However, if the same system is used to determine access to an essential service or make eligibility decisions, it may require deeper assessment and potentially higher compliance controls.

2. AI Risk Classification

Risk classification determines the level of compliance effort required for each AI system. Under the EU AI Act, AI systems are generally assessed through a risk-based model that includes unacceptable risk, high risk, limited risk, and minimal risk. Each category creates different obligations, from prohibition to transparency controls or voluntary governance practices.

2.1 Categorize AI Systems

Classify each AI system based on its purpose, users, context, and potential impact. The same type of AI technology may fall into different risk categories depending on how it is used. For example, a text summarization model used to summarize public marketing content is likely lower risk than an AI system used to screen job applicants or assess eligibility for financial services.

- **Unacceptable Risk:** AI practices that are prohibited because they create unacceptable risks to individuals or society. Examples may include manipulative AI practices, certain forms of social scoring, or prohibited biometric and surveillance-related uses.

- **High Risk:** AI systems that may significantly affect health, safety, or fundamental rights. Examples may include AI used in recruitment, employee evaluation, creditworthiness assessment, education access, critical infrastructure, law enforcement, migration, or access to essential services.

- **Limited Risk:** AI systems that are generally permitted but may require transparency measures. Examples include chatbots, AI-generated content tools, or systems where users should be informed that they are interacting with AI.
- **Minimal Risk:** AI systems with low expected impact that generally do not trigger mandatory EU AI Act obligations. Examples may include spam filters, basic recommendation tools, or internal productivity aids, provided they do not affect regulated or rights-sensitive decisions.

Practical classification test: Ask whether the AI system can materially affect a person's job opportunity, access to education, access to credit, access to essential services, safety, legal position, or fundamental rights. If the answer is yes, the system should be escalated for legal, compliance, and risk review before deployment or continued use.

2.2 Determine Compliance Obligations

Once each AI system is classified, identify the regulatory obligations that apply to the relevant risk category. High-risk AI systems should be prioritized because they require more extensive governance, documentation, monitoring, human oversight, risk management, and conformity-related controls. Limited-risk systems may mainly require transparency measures, while minimal-risk systems should still be managed through internal AI governance policies and voluntary good practices.

- **For unacceptable-risk systems:** Stop development, procurement, or deployment unless legal review confirms the system is outside the prohibited scope.
- **For high-risk systems:** Prioritize risk management, technical documentation, data governance, logging, transparency, human oversight, accuracy, robustness, cybersecurity, conformity assessment, and post-market monitoring activities.
- **For limited-risk systems:** Implement transparency controls, such as informing users when they are interacting with AI or when content is AI-generated where applicable.
- **For minimal-risk systems:** Maintain basic inventory records, acceptable-use rules, user guidance, and periodic review to ensure the system does not evolve into a higher-risk use case.
- **For third-party AI systems:** Review vendor documentation, contractual responsibilities, instructions for use, audit rights, incident reporting processes, and evidence of EU AI Act alignment.

Example: If a finance department uses an AI model to help assess creditworthiness, the organization should treat the system as potentially high risk and prioritize documentation, governance review, input data controls, human oversight, and monitoring. If a marketing team uses an AI writing assistant to draft campaign copy, the system may be limited or minimal risk depending on whether customers are exposed to AI-generated content and whether transparency obligations apply.

3. Data Governance

Data governance is a central control area for AI compliance because AI system performance, fairness, reliability, and auditability depend heavily on the quality and management of data. For high-risk AI systems, organizations should be able to explain how training, validation, and testing datasets were selected, prepared, reviewed, and monitored. Strong data governance also helps demonstrate that the organization has considered bias, data gaps, representativeness, retention, access, and lawful usage.

3.1 Review Data Quality

Verify that training, validation, and testing datasets are accurate, complete, relevant, and representative for the intended purpose of the AI system. A dataset that works well in one context may not be suitable in another. For example, a model trained on data from one geography, language, customer segment, or employment population may perform poorly or unfairly when deployed in a different environment.

- **Check accuracy:** Confirm that source records are correct, current, and not materially corrupted by duplicates, missing values, outdated labels, or inconsistent formats.
- **Check completeness:** Identify whether important fields, categories, populations, regions, languages, or usage scenarios are missing from the dataset.
- **Check representativeness:** Assess whether the dataset reflects the real-world context in which the AI system will be used.

- **Assess bias risk:** Review whether the data may create unfair outcomes for individuals or groups, especially where outputs could affect employment, education, credit, public services, safety, or access to opportunities.
- **Review data preparation steps:** Document cleaning, annotation, labelling, enrichment, aggregation, sampling, and transformation activities.
- **Monitor data drift:** Establish a process to detect when live operating data begins to differ from the data used to train or validate the model.

Example: If an organization uses an AI model to support recruitment screening, the compliance team should check whether historical hiring data reflects past bias or underrepresentation. If the dataset mainly includes applicants from a narrow set of universities, locations, or demographic profiles, the model may learn patterns that unfairly disadvantage qualified candidates from other backgrounds.

3.2 Document Data Sources

Maintain clear records of data origin, lawful usage, access permissions, and retention rules. Data source documentation should show where the data came from, why it was collected, how it was prepared, and whether it can be used for the AI system's intended purpose. This is especially important where personal data, sensitive data, employee data, customer data, or third-party datasets are involved.

- **Record data origin:** Document whether the data came from internal systems, customer interactions, public sources, vendor datasets, synthetic data, open-source repositories, or third-party platforms.
- **Define permitted use:** Confirm that the organization has a lawful and policy-approved basis for using the data in training, testing, monitoring, or operating the AI system.
- **Maintain source lineage:** Link datasets to system records so reviewers can trace how data moved from source collection to model development or deployment.
- **Define retention rules:** Specify how long raw data, transformed data, model training data, test data, logs, prompts, outputs, and evaluation records will be retained.
- **Set access controls:** Limit access to datasets based on job role, business need, security classification, and privacy requirements.
- **Document restrictions:** Capture contractual, regulatory, privacy, intellectual property, or geographic restrictions that may limit how data can be used.

Example: If a customer service AI assistant is trained using historical support tickets, the organization should document whether those tickets contain personal data, whether the original collection notice allowed this use, who may access the dataset, how long support conversations are retained, and whether vendor processing is covered by contract.

4. Transparency Requirements

Transparency requirements help ensure that individuals understand when they are interacting with AI or receiving AI-generated or AI-manipulated content. These controls reduce the risk of deception, support informed decision-making, and strengthen trust in AI-enabled services. Transparency should be designed into the user experience rather than hidden in long legal terms or technical documentation.

4.1 AI Disclosure

Inform users when they are interacting with an AI system, especially where the AI system directly communicates with people through chat, voice, email, virtual assistant interfaces, automated support workflows, or agentic actions. AI-generated content should also be appropriately identified when required, particularly where synthetic text, audio, images, or video may affect user understanding or public information.

- **Use clear disclosure language:** Tell users upfront when they are interacting with AI, such as "You are chatting with an AI assistant."
- **Place disclosures where users will see them:** Avoid relying only on terms and conditions, footnotes, or privacy notices that users may not read before interacting with the system.
- **Identify AI-generated content:** Label or mark AI-generated text, images, audio, or video where required by the use case.

- **Address deepfake risk:** Ensure manipulated or synthetic media that resembles real people, events, voices, or visuals is disclosed appropriately.
- **Document disclosure decisions:** Keep records showing why a disclosure was required, how it was implemented, and who approved it.
- **Review exceptions carefully:** Do not assume that an exception applies unless legal or compliance review confirms it.

Example: A website chatbot that answers customer questions should clearly state that it is an AI assistant before or at the start of the conversation. If the chatbot escalates to a human agent, the transition should also be clear so users understand whether they are communicating with AI or a person.

4.2 User Communication

Review customer-facing AI tools for transparency compliance and update notices, disclosures, scripts, help-center content, and user journeys where necessary. Communication should be understandable, timely, and matched to the user's context. A technically correct disclosure may still be ineffective if it appears too late, is difficult to notice, or is written in language that users cannot easily understand.

- **Review customer-facing channels:** Check websites, mobile apps, call center tools, chatbots, AI email responders, automated claims systems, recommendation engines, and self-service portals.
- **Update notices and disclosures:** Refresh privacy notices, AI notices, chatbot welcome messages, product documentation, and customer support scripts.

- **Use plain language:** Explain what the AI does, what users can expect, and how they can seek help or human review where relevant.
- **Coordinate across teams:** Align Legal, Compliance, Product, Marketing, Customer Support, Privacy, Security, and UX teams before launching or changing AI disclosures.
- **Test user understanding:** Validate whether users actually notice and understand the disclosure during normal use.
- **Maintain evidence:** Keep screenshots, approval records, change logs, and version history for audit readiness.

Example: If a marketing team publishes AI-generated product descriptions or campaign images, the organization should determine whether labels, metadata, watermarking, or other disclosures are required. The review should consider the channel, audience, whether the content could mislead users, and whether the content relates to matters of public interest.

5. Human Oversight

Human oversight is a key safeguard for AI systems, especially where outputs may influence decisions about people, safety, legal rights, financial access, employment, education, or essential services. Oversight should not be treated as a symbolic approval step. It should be designed so that qualified personnel can understand the AI system's output, question it, intervene when needed, and escalate concerns before harm occurs.

5.1 Assign Accountability

Identify responsible owners for every AI system and define who is accountable for business use, technical performance, legal compliance, data governance, security, and operational monitoring. Accountability should be documented before deployment and reviewed whenever the system changes purpose, user group, data source, vendor, model version, or decision impact.

- **Assign a business owner:** Identify the department leader responsible for the AI-enabled process and its outcomes.
- **Assign a technical owner:** Identify the team responsible for model operation, integrations, monitoring, troubleshooting, and change management.
- **Assign a compliance owner:** Identify who confirms EU AI Act, privacy, security, sectoral, and internal policy requirements.

- **Define approval gates:** Require documented approval before pilot launch, production deployment, major model updates, vendor changes, and expansion to new use cases.
- **Create escalation routes:** Define when issues should be escalated to Legal, Compliance, Risk, Security, Data Protection, Human Resources, or senior management.
- **Clarify vendor accountability:** For third-party AI systems, document which obligations sit with the provider, deployer, importer, distributor, or internal business owner.

Example: If an AI tool is used to prioritize fraud alerts in a finance function, the fraud operations manager may be the business owner, the data science or IT team may be the technical owner, and Compliance or Risk may be responsible for confirming that monitoring, escalation, and audit controls are operating effectively.

5.2 Enable Human Intervention

Ensure qualified personnel can review, challenge, pause, override, or reverse AI-generated decisions where appropriate. Human intervention should be practical, timely, and supported by clear procedures. Reviewers must have enough information to understand the AI output, the relevant context, the limits of the system, and the steps available when the output appears incorrect, biased, unsafe, or inappropriate.

- **Define review authority:** Specify who can approve, reject, override, suspend, or escalate AI-supported decisions.
- **Provide reviewer guidance:** Give human reviewers instructions on interpreting AI outputs, confidence levels, risk warnings, and known limitations.
- **Establish exception handling:** Create procedures for ambiguous outputs, suspected bias, data errors, system failures, complaints, and adverse decisions.
- **Enable pause or rollback:** Ensure teams can stop use of an AI system or revert to a manual process when the system behaves unexpectedly.
- **Train oversight personnel:** Provide role-specific training so reviewers understand both the business process and the AI system's limitations.
- **Track intervention decisions:** Record when humans reviewed, changed, rejected, or escalated AI outputs so the organization can demonstrate effective oversight.

Example: If an AI system recommends whether a customer should be approved for a loan, a trained reviewer should be able to examine the relevant inputs, understand the system's recommendation, request additional information, override the recommendation where justified, and record the reason for the final decision.

6. Documentation & Audit Readiness

Documentation and audit readiness help demonstrate that AI governance controls are not only designed but actually operating. For high-risk AI systems, documentation should be clear, current, and detailed enough for internal reviewers, external auditors, regulators, competent authorities, and notified bodies to understand the system, assess compliance, and trace important governance decisions across the AI lifecycle.

6.1 Maintain Documentation

Keep AI system documentation up to date throughout the lifecycle of the system. Documentation should explain what the system does, how it is intended to be used, what data it relies on, how it was tested, what controls are in place, and what governance decisions were made. It should also capture changes over time so reviewers can understand which version was in use at a particular date.

- **Maintain system descriptions:** Document the AI system's purpose, scope, users, operating environment, outputs, limitations, and dependencies.
- **Record model versions:** Track model names, version numbers, release dates, configuration changes, prompts, model cards, and deployment environments.
- **Record testing results:** Keep evidence of validation, performance testing, bias testing, robustness checks, security testing, and user acceptance testing.

- **Capture governance decisions:** Record approvals, risk assessments, legal reviews, privacy assessments, security reviews, vendor assessments, and exception approvals.
- **Update after changes:** Refresh documentation when models are retrained, data sources change, system behavior changes, vendors update functionality, or the use case expands.
- **Link documentation to controls:** Connect each compliance requirement to supporting evidence, such as policies, logs, test reports, approvals, training records, and monitoring dashboards.

Example: For an AI system used in employee performance support or workforce planning, documentation should describe the intended purpose, the data used, whether outputs are advisory or decision-making, who can access results, how bias was assessed, and what human review steps are required before any employment-related action is taken.

6.2 Prepare Audit Evidence

Maintain audit logs and compliance records that can support regulatory reviews, internal audits, vendor assurance reviews, customer due diligence, and management reporting. Audit evidence should be organized, version-controlled, access-controlled, and retained according to legal, regulatory, contractual, and internal policy requirements.

- **Maintain automatic logs:** Ensure relevant AI systems can generate logs showing important events, usage, outputs, errors, human reviews, and operational incidents.
- **Store compliance evidence:** Keep risk assessments, test results, data governance records, transparency notices, human oversight records, vendor documents, and approval records.
- **Protect evidence integrity:** Use access controls, retention schedules, version history, tamper protection, and clear ownership for audit repositories.
- **Prepare audit trails:** Ensure auditors can trace a requirement from policy to control design, implementation evidence, operating evidence, and management review.
- **Review evidence periodically:** Schedule evidence refresh cycles so records remain current and do not become stale before an audit or regulatory inquiry.
- **Document corrective actions:** Record issues found during monitoring, audit testing, incidents, complaints, or regulatory reviews, along with owners, due dates, remediation actions, and closure evidence.

Example: If an auditor asks how a high-risk AI system was approved for production, the organization should be able to provide the risk assessment, technical documentation, data quality review, testing results, human oversight design, approval record, deployment date, monitoring logs, and any post-deployment issue records.

7. Security & Risk Controls

Security and risk controls protect AI systems, sensitive data, model assets, prompts, outputs, integrations, and downstream business processes from misuse, unauthorized access, manipulation, and disruption. For high-risk AI systems, security should cover both traditional cybersecurity risks and AI-specific threats such as data poisoning, model poisoning, prompt injection, adversarial inputs, model leakage, unauthorized model changes, and confidentiality attacks.

7.1 Review Security Measures

Protect AI systems and sensitive data from unauthorized access throughout the AI lifecycle. Security measures should apply from data collection and model development through deployment, monitoring, incident response, and retirement. The organization should treat AI systems as critical technology assets where compromise could affect confidentiality, integrity, availability, business operations, regulatory compliance, or individual rights.

- **Apply access controls:** Restrict access to AI models, datasets, logs, prompts, outputs, source code, APIs, and administrative consoles based on least privilege and role-based access.
- **Protect sensitive data:** Use encryption, masking, tokenization, secure storage, and approved transfer channels for personal data, confidential business data, regulated records, and proprietary model assets.

- **Secure AI integrations:** Review APIs, plugins, automation workflows, agent permissions, identity controls, and connections to enterprise systems.
- **Monitor AI-specific threats:** Implement controls to detect and respond to prompt injection, adversarial examples, data poisoning, model poisoning, unauthorized model changes, and abnormal output behavior.
- **Perform security testing:** Include vulnerability testing, penetration testing, red teaming, abuse-case testing, and secure configuration review for AI-enabled applications.
- **Define incident response:** Establish procedures for AI-related incidents, including data leakage, unsafe outputs, system compromise, model manipulation, unexpected automation, and regulatory reporting triggers.

Example: If a customer support AI assistant connects to order history, billing systems, and case records, the security review should verify that the assistant cannot expose customer data to unauthorized users, execute actions beyond approved permissions, or retrieve restricted information through prompt manipulation.

7.2 Align with Standards

Review AI governance practices against ISO/IEC 42001 and ISO/IEC 27001 requirements. ISO/IEC 42001 supports an Artificial Intelligence Management System focused on responsible AI governance, risk management, transparency, accountability, lifecycle controls, and continual improvement. ISO/IEC 27001 supports an Information

Security Management System focused on protecting information through confidentiality, integrity, availability, access control, risk treatment, and security governance. Together, these standards provide a practical structure for aligning AI governance with broader enterprise security and risk management.

- **Map AI controls to ISO/IEC 42001:** Review policies, roles, AI risk assessments, impact assessments, lifecycle controls, monitoring, performance evaluation, and corrective action processes.
- **Map security controls to ISO/IEC 27001:** Review access management, asset inventory, supplier security, incident management, business continuity, secure development, logging, and vulnerability management.
- **Identify security gaps:** Compare current AI practices against approved policies, regulatory obligations, security baselines, and risk appetite.
- **Address AI-specific risks:** Include controls for bias, explainability, model drift, prompt security, model access, training data integrity, and human oversight.
- **Coordinate with enterprise risk teams:** Align AI risk registers with cybersecurity, privacy, operational resilience, third-party risk, and compliance risk registers.
- **Track remediation:** Assign owners, due dates, priorities, evidence requirements, and management reporting for every identified gap.

Example: An organization may already have ISO 27001 controls for access management, vulnerability management, and incident response. To support AI compliance, it should extend those controls to AI-specific assets such as model repositories, prompt libraries, training datasets, evaluation records, and AI monitoring dashboards, while using ISO/IEC 42001 to structure AI governance roles, lifecycle controls, and continual improvement.

8. Continuous Monitoring

Continuous monitoring ensures that AI systems remain accurate, reliable, secure, compliant, and aligned with their intended purpose after deployment. AI systems can change in performance over time because data patterns shift, user behavior changes, vendors update models, new risks emerge, or the system is used in ways that were not originally anticipated. Monitoring should therefore cover technical performance, operational behavior, data changes, user feedback, incidents, and compliance status.

8.1 Monitor AI Performance

Track model performance, accuracy, robustness, drift, and output quality against defined metrics. Monitoring should be proportional to risk: a low-risk internal productivity tool may need light-touch monitoring, while a high-risk AI system should have formal metrics, alerts, thresholds, review cadence, and documented remediation steps.

- **Define performance metrics:** Track accuracy, error rates, false positives, false negatives, confidence levels, latency, uptime, user feedback, and business outcome indicators.
- **Monitor data drift:** Compare live input data with training, validation, and testing datasets to identify changes in patterns, quality, completeness, or representativeness.

- **Monitor model behavior:** Review output consistency, hallucination risk, unsafe outputs, unexpected recommendations, automated actions, and failure patterns.
- **Set thresholds and alerts:** Define when performance degradation, abnormal behavior, bias indicators, or security signals require investigation or escalation.
- **Review vendor changes:** Track changes made by third-party providers, including model updates, feature releases, configuration changes, and revised instructions for use.
- **Record monitoring outcomes:** Maintain logs, dashboards, review notes, incident records, and corrective actions as audit evidence.

Example: A fraud detection model may initially perform well but become less accurate as fraud patterns change. Monitoring should track false positives, false negatives, investigation outcomes, data drift, and analyst overrides so the organization can decide whether retraining, threshold adjustment, additional controls, or temporary manual review is needed.

8.2 Conduct Periodic Reviews

Schedule regular governance and compliance reviews to confirm that AI systems remain within their approved scope, risk classification, operating assumptions, and regulatory obligations. Reviews should also occur after significant model changes, new data sources, vendor updates, new geographies, new user groups, major incidents, customer complaints, or changes in law or organizational policy.

- **Set review frequency:** Define monthly, quarterly, semi-annual, or annual review cycles based on system risk, business criticality, and regulatory exposure.
- **Review risk classification:** Confirm whether the system remains unacceptable, high, limited, or minimal risk based on its current use and impact.
- **Update documentation:** Refresh system descriptions, model versions, data sources, test results, user instructions, vendor records, monitoring reports, and governance approvals after significant changes.
- **Review incidents and complaints:** Analyze user complaints, adverse outcomes, security incidents, bias concerns, exception cases, and human intervention records.
- **Evaluate control effectiveness:** Test whether oversight, transparency, security, data governance, and documentation controls continue to operate as designed.
- **Report to governance forums:** Summarize review outcomes, open risks, remediation progress, and decisions for AI governance committees, risk committees, or senior management.

Example: If an HR chatbot originally answered policy questions but is later configured to recommend disciplinary actions or evaluate employee performance data, the organization should trigger a new risk classification, legal review, data governance assessment, transparency review, and human oversight design before the expanded use continues.

9. Compliance Readiness Review

A compliance readiness review brings together the results of inventory, classification, data governance, transparency, human oversight, documentation, security, and monitoring activities. The purpose is to confirm whether the organization is ready to demonstrate compliance, identify unresolved gaps, and prioritize remediation before regulatory review, customer due diligence, internal audit, or board-level reporting.

9.1 Verify Key Requirements

Use the readiness checklist below as a final control gate before approving an AI system for continued use, production deployment, procurement renewal, or expanded business use. Each item should be supported by evidence, not only by verbal confirmation. Where an item is incomplete, record the gap, owner, due date, risk rating, and interim control.

- ☐ AI systems inventoried
- ☐ Risk classifications completed
- ☐ Data governance reviewed
- ☐ Bias assessments conducted
- ☐ Transparency requirements implemented
- ☐ Human oversight established

- ☐ Documentation updated
- ☐ Security controls validated
- ☐ Continuous monitoring in place
- ☐ Audit evidence available

Example: Before renewing a vendor contract for an AI-enabled recruitment tool, the organization should verify that the tool is inventoried, classified, supported by vendor documentation, reviewed for bias risk, covered by human oversight procedures, protected by security controls, and supported by evidence that can be presented during audit or regulatory inquiry.

9.2 Address Compliance Gaps

Prioritize remediation activities based on risk level, regulatory exposure, business criticality, affected users, and the seriousness of the control gap. High-risk AI systems and systems affecting individuals' rights or access to important services should receive priority attention. Gap remediation should be tracked through a formal action plan rather than handled informally.

- **Rank gaps by severity:** Separate critical, high, medium, and low-risk gaps based on potential harm, legal exposure, control weakness, and likelihood of occurrence.

- **Assign accountable owners:** Name the business, technical, compliance, security, privacy, or vendor owner responsible for each remediation activity.
- **Set timelines:** Define target completion dates, interim milestones, escalation triggers, and management reporting cadence.
- **Define interim controls:** Where remediation will take time, apply temporary safeguards such as manual review, restricted access, additional monitoring, reduced automation, or suspension of specific features.
- **Track closure evidence:** Require documentation showing that the gap was fixed, tested, approved, and incorporated into ongoing governance.
- **Escalate unresolved risks:** Report overdue or high-impact gaps to the AI governance committee, risk committee, executive sponsor, or board-level forum where appropriate.

Example: If a high-risk AI system has incomplete logging, weak human override procedures, or undocumented training data sources, the organization should not treat the issue as a minor documentation task. It should assign owners, apply interim controls, set deadlines, and confirm closure through testing and evidence review.

10. Next Steps

After completing the checklist, organizations should move from one-time compliance review to an ongoing AI governance operating model. The EU AI Act should be treated as part of a broader responsible AI, risk management, information security, privacy, legal, and operational resilience program. The next steps below help convert assessment findings into a sustainable roadmap.

10.1 Build Your AI Governance Roadmap

Develop an AI governance strategy aligned with business objectives, risk appetite, regulatory obligations, technology strategy, and customer trust goals. The roadmap should define how AI systems are approved, monitored, documented, reviewed, and retired. It should also clarify decision rights, funding needs, timelines, and governance forums.

- **Define governance structure:** Establish an AI governance committee or working group with representatives from Legal, Compliance, Risk, Security, Privacy, IT, Data Science, HR, Procurement, and business functions.
- **Create approval workflows:** Define how AI systems move from idea to pilot, production, monitoring, and retirement.
- **Standardize templates:** Use common templates for AI inventory, risk classification, impact assessment, vendor review, data governance, human oversight, testing, and audit evidence.

- **Integrate with existing programs:** Align AI governance with enterprise risk management, information security, privacy, model risk management, vendor risk management, internal audit, and compliance monitoring.
- **Prioritize roadmap phases:** Sequence actions into immediate, short-term, medium-term, and long-term workstreams.
- **Report progress:** Provide regular updates to senior management or governance forums on AI risks, remediation status, compliance readiness, and upcoming regulatory developments.

Example roadmap: In the first 30 days, complete the AI inventory and identify high-risk systems. In 60 days, complete risk classification, ownership mapping, and gap assessment. In 90 days, implement priority controls for data governance, transparency, human oversight, documentation, and monitoring. After that, transition to quarterly governance reviews and annual audit readiness testing.

10.2 Strengthen Organizational Readiness

Train relevant teams on AI governance and regulatory requirements so that compliance is embedded into daily decision-making. Organizational readiness depends on more than policies; it requires people to understand their responsibilities, recognize AI risk, use approved processes, and know when to escalate issues. Training should be role-based and refreshed as AI systems, regulations, and internal policies evolve.

- **Train business teams:** Help product owners, process owners, HR, Finance, Marketing, Operations, and Customer Support identify AI use cases, risk triggers, and required approvals.
- **Train technical teams:** Cover data governance, model testing, monitoring, cybersecurity, documentation, prompt security, and change management.
- **Train oversight personnel:** Ensure human reviewers understand when and how to question, override, pause, or escalate AI outputs.
- **Train procurement and vendor teams:** Build capability to review vendor AI documentation, contractual responsibilities, audit rights, incident reporting, and change notifications.
- **Refresh policies and procedures:** Update acceptable-use policies, AI development standards, vendor onboarding procedures, incident response plans, and audit evidence requirements.
- **Review compliance posture regularly:** Reassess the organization's AI governance program as AI regulations, standards, regulator guidance, business use cases, and technology capabilities continue to evolve.

Example: A company deploying generative AI assistants across multiple departments may create separate training tracks for general users, prompt designers, system administrators, legal reviewers, security teams, and senior decision-makers. This ensures that each group understands the risks and responsibilities relevant to its role.

Checklist Summary

- Create and maintain a centralized AI inventory.
- Assign business, technical, and compliance owners for each AI system.
- Map AI usage by department, process, user group, and decision impact.
- Identify whether each AI system is internally developed, vendor-provided, open-source-based, or embedded in another platform.
- Classify each AI system as unacceptable risk, high risk, limited risk, or minimal risk.
- Escalate high-risk and borderline systems for legal, compliance, privacy, security, and business review.
- Prioritize high-risk systems for documentation, oversight, monitoring, and control implementation.
- Review vendor-provided AI systems for contractual accountability, transparency, and evidence of compliance readiness.
- Review training, validation, and testing datasets for accuracy, completeness, representativeness, and bias risk.
- Document data sources, data lineage, permitted usage, retention periods, and access controls.

- Inform users when they are interacting with AI systems where required.
- Label or identify AI-generated or AI-manipulated content where transparency obligations apply.
- Review customer-facing AI tools, notices, and disclosures for transparency compliance.
- Maintain screenshots, approvals, logs, and version records as evidence of transparency implementation.
- Assign business, technical, and compliance accountability for every AI system.
- Define approval gates, escalation routes, and exception-handling procedures.
- Ensure qualified personnel can review, challenge, override, pause, or escalate AI-generated decisions where appropriate.
- Train oversight personnel on AI limitations, decision context, and intervention procedures.
- Maintain current AI system documentation, including model versions, testing results, control evidence, and governance decisions.
- Store audit logs and compliance records in controlled repositories with retention, access, and integrity safeguards.
- Prepare audit trails that connect regulatory requirements to policies, controls, evidence, monitoring, and corrective actions.

- Protect AI systems, datasets, logs, prompts, outputs, APIs, and administrative interfaces from unauthorized access.
- Implement cybersecurity controls for AI-specific threats such as prompt injection, data poisoning, model poisoning, adversarial inputs, and confidentiality attacks.
- Align AI governance with ISO/IEC 42001 and information security controls with ISO/IEC 27001.
- Identify security and governance gaps, assign remediation owners, and track closure evidence.
- Monitor AI performance, accuracy, robustness, data drift, model behavior, user feedback, and vendor changes.
- Set thresholds and escalation triggers for performance degradation, abnormal outputs, bias indicators, or security signals.
- Conduct periodic governance and compliance reviews based on system risk and business criticality.
- Update documentation after significant model changes, new data sources, vendor updates, incidents, or expanded use cases.
- Complete a final compliance readiness review before deployment, renewal, expansion, or audit.

- Verify that inventory, risk classification, data governance, bias assessment, transparency, human oversight, documentation, security, monitoring, and audit evidence are complete.
- Prioritize compliance gaps by severity, assign accountable owners, define timelines, and track closure evidence.
- Build an AI governance roadmap aligned with business objectives, risk appetite, regulatory obligations, and customer trust goals.
- Train relevant teams on AI governance, AI literacy, security, human oversight, vendor responsibilities, and regulatory requirements.
- Schedule periodic reclassification reviews when AI systems change purpose, data, users, geography, or decision impact.

GSDC
Global Skill Development Council

CERTIFIED ISO 27001:2022 LEAD AUDITOR

ISO 27001 LEAD AUDITOR CERTIFICATION IS BASED ON INFORMATION SECURITY MANAGEMENT SYSTEM (ISMS) AND GLOBAL COMPLIANCE STANDARDS

ABOUT GSDC CERTIFICATION

EBOOK

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

CREATED BY EXPERTS

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

EBOOK

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

LEARNING MATERIALS

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

LEARNING OBJECTIVE

- Verify adherence to legal and regulatory requirements
- Provide recommendations for enhancing ISMS
- Ensure continuous improvement of security practices
- Foster a culture of risk management awareness

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

Enroll Now

www.gsdcouncil.org