Preparation checklist

AI Audit Snapshot · for a CISO or security lead · 87 questions, about 72 minutes with the answers at hand

How to use this. These are the questions you will be asked, in the order you will see them. Tick each one you can answer today; for the rest, note who in the company has the answer. You do not need to finish in one sitting: leave the questionnaire at any point and the resume link in your email brings you back with everything saved. "Don't know" is always an allowed answer and is treated as a finding about visibility, not as a mark against you.

Context and identity

Have at hand: Your industry code, headcount, revenue band, where you operate, which laws and standards you work under, and who owns AI decisions today.

- ☐ How long have you been in this role?
- ☐ Which AI-relevant laws, regulations, or frameworks does your company operate under? Check every one that applies (or would apply if a regulator asked).
- ☐ How long has your company been using any AI tools?
- ☐ Was your AI adoption strategy formally approved by leadership before tools were deployed?
- ☐ Why are you completing this assessment? (select all that apply)
- ☐ Voluntary AI, data, and security standards your company aligns with or claims compliance to. Includes standards required by customer contracts or insurance policies.
- ☐ Governance and fiduciary frameworks that bear on your board's or executive team's AI oversight obligations.
- ☐ Contractual and supply chain expectations that require AI-related evidence from your company.
- ☐ Across the regulatory, voluntary, governance, and contractual items you selected in the previous four questions, what is the dominant cadence at which your company must produce evidence?

AI tool inventory and discovery

Have at hand: A list of every AI tool in use, sanctioned or not: name, who uses it, what data it touches, who pays for it. Ask around; the unsanctioned ones matter most.

- ☐ Which AI tools does your company currently use? Check every tool anyone at the company uses.
- ☐ Does a single document exist that lists every AI tool currently in active use across the company?
- ☐ What does that inventory include? (select all that apply)

- ☐ When was the inventory last updated?
- ☐ Who maintains the inventory?
- ☐ Has anyone attempted to create an AI tool inventory?
- ☐ Why hasn't an inventory been built? (select all that apply)
- ☐ If a regulator or acquirer asked for a complete AI tool inventory today, how long would it take to produce?
- ☐ Have you been asked for an AI tool inventory in the past 3 months by anyone external (auditor, insurer, customer, board)?
- ☐ How many AI tools does leadership believe are in use across the company?
- ☐ How confident is that number?
- ☐ Have any AI tools been adopted at the team or department level without central IT or leadership approval?
- ☐ To what extent has the company identified vendor-enabled AI features in its existing software stack (in the past 3 months)?
- ☐ Has the company conducted endpoint or network discovery to identify AI tool usage?
- ☐ For the top 3 AI tools by spend or strategic importance in your company, indicate which of the following apply:
  - ☐ Tool 1
  - ☐ Tool 2
  - ☐ Tool 3
- ☐ Which categories of vendor-enabled AI features are confirmed active in your environment? (select all that apply)
- ☐ How many AI tools have been added in the past 3 months?
- ☐ How many AI tools have been retired in the past 3 months?
- ☐ Were any tools retired because they failed to deliver value, or because of risk/compliance concerns?
- ☐ Is there a defined approval process required before a new AI tool is adopted?

Cost mapping

Have at hand: Last 12 months of AI spend: subscriptions, API bills, vendor contracts, staff time. Your finance system or card statements are the source.

- ☐ Which of the following hidden AI costs has the company quantified or documented (in the past 3 months)? (select all that apply)

Performance measurement

Have at hand: Any targets set for AI tools, any measurement of results, any tool retired for missing its target.

- ☐ For the company's top 3 AI use cases by spend or strategic importance, indicate which apply (past 3 months):
  - ☐ Use case 1
  - ☐ Use case 2
  - ☐ Use case 3
- ☐ Has any AI tool been continued despite consistently missing its target metric?
- ☐ Has any AI tool been killed because it missed its target metric?
- ☐ How often are AI errors caught before reaching customers or final decisions?
- ☐ Is there a defined review process for high-stakes AI outputs (customer communications, financial decisions, hiring)?
- ☐ Has any AI error reached a customer or external party in the past 3 months?
- ☐ Has any AI error been disclosed to a regulator or auditor?
- ☐ Have you measured customer or employee satisfaction with AI-touched processes?
- ☐ For AI systems involved in hiring, promotion, lending, credit, or other regulated decisions, has a documented bias audit or fair lending review been conducted in the past 12 months?

Risk exposure

Have at hand: Vendor contracts and data-handling terms, where customer or regulated data goes, any incident or near miss, cyber insurance terms, counsel's involvement.

- ☐ Has sensitive company information been entered into any external AI tool?
- ☐ What categories of sensitive information? (select all that apply)
- ☐ Have AI vendor data-handling terms been reviewed by legal counsel?
- ☐ Have any AI tools been confirmed to retain or train on submitted data?
- ☐ Have data subjects (customers, employees) been notified about AI processing where required by law?
- ☐ Have you mapped which AI tools touch regulated data?
- ☐ Is there a documented data classification policy in place?
- ☐ Are AI tools constrained by data classification?
- ☐ Has AI been involved in any of the following high-stakes decisions or outputs in the past 3 months? (select all that apply)
- ☐ Are AI tools generating content that could be misattributed to a human?
- ☐ Are any AI tools embedded in customer-facing communications?
- ☐ Have any AI vendors disclosed a security incident affecting your data in the past 3 months?
- ☐ Have you confirmed AI vendor SOC 2 or ISO compliance?

- ☐ If AI is involved in hiring, promotion, or termination decisions, has a bias audit been conducted in the past 12 months?
- ☐ Are any AI tools operating in an autonomous or agentic mode — taking actions, executing transactions, or modifying systems without per-action human approval?
- ☐ For AI systems operating autonomously, has the company defined the approved scope of autonomous action (what the AI is allowed to do unsupervised)?
- ☐ For AI systems operating autonomously, are there documented circuit breakers — defined conditions under which the AI must stop and escalate to a human?
- ☐ For AI systems operating autonomously, are there gating controls on the external systems the AI can access (CRM, finance system, customer-facing channels)?
- ☐ For AI vendors handling regulated or sensitive data, which contractual provisions exist? (select all that apply)
- ☐ Does your company publish a customer-facing list of AI sub-processors (publicly available or available to customers on request) — the customer-facing AI Third-Party Governance Statement?
- ☐ Does your company use a standard AI-specific procurement rider or contract addendum when entering new AI vendor agreements?
- ☐ For AI use cases affecting fundamental rights (employment, credit, healthcare, education, housing, public benefits), has human override capability been designed into each use case?

Governance gaps

Have at hand: Board or executive minutes mentioning AI, any AI policy, the committee or person accountable, training records, how exceptions are handled.

- ☐ Is there a published AI usage policy?
- ☐ Is there a documented process for evaluating new AI tools before adoption?
- ☐ Is there a documented incident response process for AI-related failures?
- ☐ Is there a training program for employees on appropriate AI use?
- ☐ What are the top three barriers to better AI outcomes in the company? (rank top 3)
- ☐ Has the company assessed its AI use cases against ISO/IEC 42001 (the AI management system standard)?
- ☐ Has the company assessed its AI use cases against the NIST AI Risk Management Framework?
- ☐ For AI systems with EU customers, employees, or partners, have they been classified against the EU AI Act risk tiers?
- ☐ For AI systems with EU AI Act exposure, has provider vs. deployer position been identified for each system?
- ☐ Are per-use-case governance records (dossier, profile, or equivalent) maintained for the highest-impact AI systems?

- ☐ For your highest-impact AI use cases (per T1-F-019 classification), does each have a system-level documentation package including: purpose, training data sources, model version, evaluation results, known limitations, and human-oversight controls?
- ☐ For AI systems handling regulated, customer, or proprietary data, is data lineage documented (where data came from, how it was prepared, how it flows into training or inference, and where outputs go)?
- ☐ If a customer, regulator, or acquirer asked which of your AI systems are subject to sector-specific regulation (HIPAA, GLBA, FERPA, DFARS, EU AI Act high-risk, NYC LL144, etc.), could you provide a current mapping within 5 business days?
- ☐ If a regulator-notifiable AI incident occurred (e.g., NYDFS Part 500, DFARS 252.204-7012, GDPR Article 33 — 72-hour windows; DORA — 4-hour initial; NIS2 — 24-hour early warning, 72-hour notification, 1-month final report), how confident is the company in meeting the shortest applicable notification deadline?
- ☐ Has your company adopted an AI-specific operational risk taxonomy (e.g., the V1 Ch 7 AI Operational Risk Categories or equivalent), with each category mapped to entries in the corporate risk register?
- ☐ Does your company maintain a documented AI internal audit log capturing scope, findings, and evidence for each review (the AI Internal Audit Workpaper Log)?
- ☐ Does your company maintain documented minutes of management reviews specifically covering AI risk, performance, and incidents (the AI Management Review Minutes Log)?
- ☐ Does your company maintain a documented register of corrective actions taken in response to AI incidents, audit findings, or management review items (the AI Corrective Action Register)?
- ☐ During an AI-related incident, does your company have a documented communication protocol covering internal stakeholders (legal, board, executives), external regulators, and affected customers (the AI Communication Alignment Protocol)?

Outcomes, workflow and confidence

Have at hand: Which business outcomes AI was meant to move, where it changed a workflow, and how confident you are in each answer.

- ☐ Confidence that the company is getting value from its AI spend
- ☐ Confidence that the company would survive regulatory scrutiny of its AI use today
- ☐ Confidence that AI risks are understood at the leadership level
- ☐ Confidence that the company's AI program is aligned with stated business outcomes

Follow-up

Have at hand: Anything the questionnaire did not ask that you think the auditor should know. Write it down; it is read.

- ☐ Which of the following would you like to receive based on these findings? (select all that apply)
- ☐ What's your preferred follow-up method?
- ☐ Anything we should know that the questionnaire didn't ask? (optional, max 1000 characters)

Questions about this checklist or the audit: srj@srjconsultingservices.com, 415-413-7772. A person answers.