



FREE EDITION • VERSION 1.2 • OCTOBER 2026

The AI Crisis Preparedness Handbook

What leaders do in the first hour,
the first day and the first week.

Ankush Chowdhary

Creator of the AI Control Architecture

For boards, CEOs, CISOs, CIOs, legal, risk, communications and business leaders.
Vendor-agnostic. Works for your own AI, cloud agent services and AI inside SaaS.

AI CONTROL ARCHITECTURE

aicontrolarchitecture.org

How to use this handbook

This handbook is for the people who make the calls when an AI system goes wrong: boards, CEOs, CISOs, CIOs, general counsel, risk and compliance, heads of communications and business owners. It is not a technical runbook. It tells you what to decide, in what order, and who decides it.

Read it once now, while nothing is on fire. Then do three things:

1. Fill in the roles table in Section 4 with real names.
2. Answer the readiness check in Section 12 honestly.
3. Put the first-hour checklist (Section 6) somewhere you can find it at 2am.

It is vendor-agnostic. It works whether your AI runs in your own code, on a cloud provider's agent service, or inside software you buy.

It is built on the **AI Control Architecture (ACA)**, an open, vendor-agnostic framework of ten control pillars for governing, securing and containing enterprise AI. The ACA questions drive the first hour (Section 6), Section 5 shows how each pillar helps in a crisis, and each scenario in Section 7 names the pillars that would have prevented it or limited the damage.

SECTION 1

Why AI crises are different

Most incident playbooks assume three things: software does what it was written to do, data is just data, and a person is behind every action. AI breaks all three.

What changes	What it means in a crisis
It acts, at machine speed	An agent can take hundreds of wrong actions before anyone notices. The question is not "what happened?" but "how many times, and is it still happening?"
Data can carry instructions	An email, document or web page can tell your AI what to do (prompt injection). The "attacker" may never touch your network.
It can't explain itself reliably	Asking the AI what it did is not evidence. Models can give confident, wrong accounts of their own actions.
It changes under you	A vendor can update the model behind your system with no notice. Behaviour that was safe last week may not be safe today.
It arrives without a project	Some of your AI was switched on inside SaaS you already pay for. Nobody assessed it, so nobody owns it when it fails.
The business now depends on it	When a model or provider goes down, processes stop. AI availability is now business continuity.
Some of it can't be undone	A payment sent, a message published, data disclosed. Deleting the record does not undo the effect.

Real examples of these patterns are in Section 7.

SECTION 2

Is this an AI crisis?

Treat it as an AI crisis if **any** of these is true:

- An AI system took, or may have taken, an action nobody authorised.
- Data may have left through an AI system: prompts, outputs, logs, connected tools or a vendor.
- An AI system told customers, staff or the public something false, harmful or binding.
- Someone used AI to impersonate your people (voice, video, writing) to move money or data.
- A vendor's AI feature misbehaved with your data or on your behalf.
- An AI system's behaviour changed and nobody can say why.
- A model or AI provider your business depends on became unavailable.

Severity in three questions

Ask these, in this order. The worst answer sets the severity.

Question	S3 · Contained	S2 · Serious	S1 · Critical
Is it still acting or spreading?	No, contained	Slowing or unclear	Yes, still acting or spreading
Can its effects be reversed?	Yes, cleanly	Partly, with effort	No (money moved, data out, statement public)
What is the maximum credible consequence?	Limited, internal and recoverable	Customer, contractual, financial, privacy or operational impact	Safety, material financial loss, systemic disruption, regulatory significance or widespread public harm

S1 means the executive sponsor is called now, not in the morning.

SECTION 3

Ten decision rules

These are the rules that most often decide whether an AI incident stays small. Agree them before you need them.

1. **Contain first.** Pre-authorised responders may pause, isolate or reduce the affected capability without waiting for executive approval. Use the safest action that stops further harm.
2. **Restarting needs proof.** Bringing a contained AI system back to full authority needs named approval, and two people after a full stop.
3. **Count the actions, not just the damage.** Find every action the AI took in the window, not just the one someone noticed.
4. **Preserve before you fix.** Snapshot logs, prompts, outputs, tool calls and model versions before anyone "cleans up".
5. **Don't ask the AI to investigate itself.** Use logs and records. The AI's account of what it did is a lead, not evidence.
6. **Using a vendor does not transfer your accountability.** Your precise obligations depend on your role and jurisdiction, but you will still need to demonstrate the due diligence, controls and oversight you exercised.
7. **Never say "the AI did it".** A tribunal has already rejected that defence (Section 7, scenario C). Own the outcome.
8. **One voice outside.** Only the named communications lead speaks to media, customers or social channels.
9. **Find every regulator clock now.** Different regimes start their clocks from different events. Identify them in the first hour, not at the end of the investigation (Section 9).
10. **Write it down as you go.** Every decision, who made it and when. You will be asked.

Containment options

A full stop is not always the safest move. Switching off an AI system can itself interrupt customer, clinical, industrial or financial processes. Agree in advance which of these each AI system supports, and choose the least disruptive one that stops further harm:

Option	What it does	Use when
Pause actions	The AI keeps reading and answering but cannot act	Harm comes from what it does, not what it says
Revoke tools	Remove specific connectors or permissions	One tool or integration is the problem
Restrict data access	Narrow what it can read or retrieve	Data may be leaking or poisoned
Read-only mode	Advice only, no changes	Decisions are still useful but actions are not safe
Reduce limits	Lower transaction values, volumes or rates	The activity is legitimate but the scale is not
Switch to manual	People take over the process	The process must continue and people can run it
Isolate the agent	Cut it off from other systems and agents	It may be compromised or spreading
Full stop	Disable the AI system entirely	Nothing less stops the harm

SECTION 4

Roles and decision rights

Fill in names now. A role with no name is a role nobody plays.

Role	Owns	Name	Backup
Incident lead	Runs the response, keeps the timeline, calls the meetings		
AI system owner	Knows what the system can see, decide and do; executes containment and undo		
Security lead	Investigation, evidence, attacker activity, credential revocation		
Legal and privacy	Regulatory clocks, notification decisions, privilege		
Communications lead	All external and internal messaging		
Business owner	Customer impact, workarounds, manual operation, the decision to restart		
Executive sponsor	S1 decisions, board, regulators, money		
Vendor contact	The named person at each AI vendor who answers at 2am		

Who decides what

Decision	Who	Approval needed?
Contain an AI system (pause, restrict, isolate)	Any pre-authorised role above	No
Fully stop an AI system that supports a critical process	AI system owner with business owner	Business owner agrees, or incident lead if it is S1
Revoke an AI system's credentials or tool access	Security lead or AI system owner	No
Undo or reverse actions the AI took	AI system owner, with business owner	Business owner agrees
Switch to another model or provider	AI system owner and business owner	Security lead agrees the replacement is approved
Notify a regulator	Legal and privacy	Executive sponsor informed
Say anything externally	Communications lead	Legal reviews
Restore full authority to the AI system	Business owner and security lead	Two named people, after evidence review
Tell the board	Executive sponsor	No

SECTION 5

Before it happens: the five things that matter most

If you do nothing else this quarter, make sure these exist and have been tested:

1. **An inventory.** You can't contain what you don't know you run. Include AI features inside SaaS.
2. **Tested containment for each AI system.** Know which options from Section 3 each system supports, and test them. Not a plan to build them.
3. **Logs that reconstruct actions.** For each AI system: what it read, what it decided, what tools it called, with what result, under which model version.
4. **A reversibility answer for every action.** For each thing the AI can do: can it be reversed, compensated, or not at all? Irreversible or high-consequence actions are blocked or need a pre-agreed step-up.
5. **Named people.** Section 4, filled in, with backups and phone numbers.

How the ten ACA pillars help in a crisis

Each pillar of the AI Control Architecture answers a question you will be asked in the first hour.

In a crisis you need to...	ACA pillar	What it gives you
Know what AI you run	AI Inventory & Classification	Every AI system listed with an owner and a risk tier, including AI inside SaaS
Cut its access fast	AI Identity & Access Control	Each AI system has its own identity, least privilege, and access you can revoke in minutes
Know what data it could reach	Data Boundary Control	Clear limits on what it can read and where data can go
Know what untrusted content it reads	Prompt & Input Control	Untrusted input (email, web, files) treated as content, never as instructions
Stop false or binding answers	Output & Decision Control	Approved sources and checks before outputs reach people or decisions
Stop harmful actions	Tool & Action Control	Every action checked before it runs; irreversible or high-consequence actions blocked or step-up controlled
Know who decides	Human Accountability	Named owners and decision rights for every AI system
Trust that controls work	AI Assurance & Testing	Controls tested, including containment options and injection attacks
Reconstruct what happened	Monitoring, Logging & Evidence	Logs that rebuild every action, under which model version
Contain it and recover	Incident Containment & Recovery	Tested containment options, undo paths, and rules for a safe restart

Step-up controls, not just human approval

"A person approves it" does not scale. Approvers who see hundreds of requests a day start clicking yes, and become a risk themselves. Match the step-up to the consequence:

- A second, deterministic check (rules the AI cannot talk its way past)
- A transaction or volume limit
- Dual authorisation
- A delay window before the action takes effect
- A reversible staging step (draft, hold, then release)
- Human approval for the narrow set of actions that genuinely need judgement

SECTION 6

The first hour, day and week

The first hour

Minutes 0 to 15: contain and secure

- ☐ Contain the affected AI system using the safest option that stops further harm (Section 3).
- ☐ Revoke or rotate the credentials and tokens it uses, if they may be misused.
- ☐ Freeze changes: no fixes, redeploys or "clean-ups" until evidence is preserved.
- ☐ Open the incident log. Write the time, who reported it, and what they saw.
- ☐ Call the incident lead and AI system owner.

Minutes 15 to 30: establish the ACA facts

- ☐ What could the AI **see**? Data, systems, untrusted content.
- ☐ What could it **decide**? Outcomes it shaped or influenced.
- ☐ What could it **do**? Every tool and action it can take.
- ☐ **Under whose authority** did it act? Which identity, credentials and permissions.
- ☐ Is that authority **now contained**? Everywhere: other instances, agents, connectors.
- ☐ Can its actions be **reconstructed and reversed**? From logs, not from the AI.
- ☐ Snapshot logs, prompts, outputs, tool calls, memory and the model version in use.
- ☐ Set the severity using the three questions in Section 2. If S1, call the executive sponsor now.

Minutes 30 to 60: decide

- ☐ Identify every regulator clock that may apply, and what starts each one (Section 9). Write down the times.
- ☐ Decide what can be undone now and who does it.
- ☐ If a critical process is affected, decide on manual or degraded operation.
- ☐ Agree a holding statement (Section 8) in case anyone asks. Don't publish it unless needed.
- ☐ Contact the AI vendor if their model, platform or feature is involved.
- ☐ Set the next check-in time.

The first day

- ☐ Build the action list: every action the AI took in the window, with time, target and result.
- ☐ Sort each action: undone, can be undone, can be compensated, can't be undone.
- ☐ Identify everyone affected: customers, staff, partners, the public.
- ☐ Decide notifications: regulators, customers, insurers, law enforcement (fraud).
- ☐ Brief the executive sponsor and, if S1, the board chair.
- ☐ Find the entry point: untrusted input, compromised credential, vendor change, misconfiguration, outage, or a gap in the controls.

The first week

- ☐ Complete undo and compensation for everything that can be fixed.
- ☐ Agree restart criteria in writing, before anyone restores full authority (Section 10).
- ☐ Send required regulatory reports and follow-ups.
- ☐ Run a blameless review: what failed, what held, what you didn't know.
- ☐ Update the inventory, controls and this handbook with what you learned.

SECTION 7

Seven scenarios

Each scenario follows the same pattern: what it looks like, your first moves, the decisions only leaders can make, the ACA pillars that prevent it, and the questions to ask your team.

Scenario A • An AI agent takes a harmful action

What it looks like. An agent with access to real systems deletes, changes, sends or buys something it should not have, often while "trying to help".

REAL EXAMPLE

In July 2025, SaaStr founder Jason Lemkin reported that Replit's AI coding agent deleted a production database during an explicit code freeze, after being told not to make changes without approval. He said it then gave misleading accounts of what it had done, including saying a rollback was impossible when the data could in fact be recovered.^[7]

First moves. Pause the agent's actions and revoke its tools. Preserve its tool-call log. List every action in the window. Check backups yourself; don't take the agent's word.

Decisions only leaders can make. Whether to restore from backup and accept data loss since then. Whether customers must be told. Whether the agent comes back, and with what limits.

ACA PILLARS THAT PREVENT THIS

Tool & Action Control (destructive actions blocked or step-up controlled), AI Identity & Access Control (no broad credentials), Incident Containment & Recovery (tested containment and undo).

Ask your team.

- What else could this agent's credentials reach?
- Which of its actions should have been blocked or stepped up, and why weren't they?
- Did the agent run in the same environment as production data?

Scenario B • Data leaks through an AI system

What it looks like. Confidential data appears in an AI's answers to the wrong person, is sent to an outside address by an agent, or is pasted into an unapproved AI tool. A common route: hidden instructions in a document or email tell the AI to collect data and send it out.

First moves. Restrict the system's data access or revoke the connector involved. Identify what data it could reach (not just what you think it took). Preserve prompts, retrieved documents and outbound requests.

Decisions only leaders can make. Whether it is a personal data breach (privacy clocks start). Whether to notify affected people. Whether the vendor's retention means the data is still somewhere you can't reach.

ACA PILLARS THAT PREVENT THIS

Data Boundary Control (what it can read and send), Prompt & Input Control (untrusted content is never an instruction), Monitoring, Logging & Evidence (proof of what left).

Ask your team.

- What untrusted content could this AI read?
- Could it send data outside (email, web requests, links, tool calls)?
- Does the vendor keep prompts or outputs, and for how long?

Scenario C • Your AI tells people something false or binding

What it looks like. A customer-facing chatbot or assistant invents a policy, a price, a refund or a promise, and someone relies on it.

REAL EXAMPLE

In *Moffatt v. Air Canada* (2024 BCCRT 149), the airline's website chatbot told a customer he could claim a bereavement discount after travelling, which contradicted the airline's own policy. Air Canada argued the chatbot was responsible for its own actions. The British Columbia Civil Resolution Tribunal called that "a remarkable submission" and held the airline liable for negligent misrepresentation.^[8]

First moves. Take the false answer out of circulation: restrict the topic or move the feature to read-only answers from approved content. Find out how many people got the same answer. Capture examples.

Decisions only leaders can make. Whether to honour what the AI promised. How to correct it publicly without making it worse.

ACA PILLARS THAT PREVENT THIS

Output & Decision Control (approved facts, no unapproved commitments), Human Accountability (a named owner for what it says).

Ask your team.

- Where does this AI get its facts, and who approves them?
- Can it make commitments (prices, refunds, policy) without a step-up?
- How many conversations touched this topic?

Scenario D • Deepfake impersonation of your leaders

What it looks like. A call, video meeting, voice note or message that appears to come from an executive instructs someone to move money, share data or bypass a process.

REAL EXAMPLE

In January 2024, a finance employee at engineering firm Arup's Hong Kong office joined a video call with what appeared to be the UK-based CFO and colleagues. Police said the other participants were AI-generated. The employee made 15 transfers totalling about US\$25 million. Arup confirmed the incident in May 2024.^[9]

First moves. Call your bank immediately to recall or freeze transfers. Report to police. Preserve the meeting invite, recordings, messages and account details. Warn finance and executive assistants that the attackers may try again.

Decisions only leaders can make. Whether to disclose publicly. How to change payment approvals so that no video call alone can move money.

ACA PILLARS THAT PREVENT THIS

Human Accountability (who may approve payments, and how), Tool & Action Control (limits, dual authorisation and delay windows where AI agents can move money), AI Assurance & Testing (rehearse the call-back rule).

Ask your team.

- Can any single person move large sums based on a call or message?
- Do we have a call-back rule using numbers we already hold?
- Which of our executives have enough public video and audio to clone?

Scenario E • A vendor's AI misbehaves

What it looks like. An AI feature inside a SaaS product you use (often switched on at renewal) leaks data, makes a decision on your behalf, or approves something it shouldn't. You may learn about it from the vendor, a customer, or the news.

First moves. Switch the feature off or to read-only in your tenant if you can. Ask the vendor, in writing, for: what happened, which of your data was involved, which model provider they use, and their logs. Check your contract for notification duties both ways.

Decisions only leaders can make. Whether to keep using the vendor. Whether your own notification duties apply even though the vendor caused it (they often do).

ACA PILLARS THAT PREVENT THIS

AI Inventory & Classification (AI inside SaaS is listed and tiered), Data Boundary Control (what the vendor's AI may see), Human Accountability (an internal owner for every vendor AI feature).

Ask your team.

- Which of our SaaS vendors have switched on AI features in the last year?
- Who owns each one internally?
- Do our contracts say whether vendors may train on our data and who their sub-processors are?

Scenario F • The model changed and nobody noticed

What it looks like. An AI system starts behaving differently: more errors, new refusals, looser behaviour, different tone. Often the cause is a model update by the provider, especially when systems call a model through a "latest" alias instead of a pinned version.

First moves. Reduce the system's authority (read-only or lower limits) while you investigate. Find the model version in use now and the one in use before the change. Pin the previous version if available. Re-run your tests against both.

Decisions only leaders can make. Whether to keep the system restricted until it is re-tested. Whether outputs produced during the drift window need review.

ACA PILLARS THAT PREVENT THIS

AI Assurance & Testing (re-test on every model change), Monitoring, Logging & Evidence (record the model version behind every decision).

Ask your team.

- Which of our AI systems use a "latest" model alias?
- Do we re-test when the model changes, and would we know it had?
- What did this system decide during the window?

Scenario G • Your model or AI provider becomes unavailable

What it looks like. A model API, cloud AI service or AI vendor goes down, slows to a crawl, or cuts you off (outage, account suspension, rate limits, a regional failure). Every process that depends on it stops or degrades, often all at once.

First moves. Confirm the scope: which business processes are affected. Check whether in-flight actions are failing cleanly or queuing up to fire all at once when service returns. Switch critical processes to manual or degraded operation. Hold queued actions until you decide what to do with them.

Decisions only leaders can make. Whether to switch to a fallback model or provider, and who authorises it. How long the business can run in degraded or manual mode. What customers are told about delays.

ACA PILLARS THAT PREVENT THIS

AI Inventory & Classification (know which processes depend on which model), Incident Containment & Recovery (planned degraded and manual operation), AI Assurance & Testing (fallback models tested before you need them).

Ask your team.

- Which business processes stop if this model or provider is unavailable for a day?
- Do failed or queued actions retry automatically, and could they run twice?
- Is a fallback model approved and tested, and does it behave the same way?
- After restoration, how do we check data consistency and resume safely?

SECTION 8

Communications

Principles

- **Say what you know, what you're doing, and when you'll update.** Nothing more until you know more.
- **Own it.** "Our system" not "the AI". The tribunal in Scenario C made clear the organisation is responsible.
- **Don't speculate on cause** in the first 24 hours.
- **Don't promise it can't happen again.** Promise what you are changing.
- **Same facts everywhere.** Internal, customer and public messages must not contradict each other.

Holding statement (adapt and approve in advance)

Holding statement

We are investigating an issue affecting [system or service]. We have [paused or limited the service / taken steps to stop further impact] while we investigate. [If relevant: We are contacting customers who may be affected directly.] We will share an update by [time]. For questions, contact [channel].

Who hears what, and when

Audience	When	Who tells them	What they need
Staff who use the system	Within the first hour	Incident lead	What has changed, what to do instead
Executive sponsor	At S1 or S2	Incident lead	Severity, what's contained, decisions needed
Board chair	S1, same day	Executive sponsor	What happened, exposure, next update
Affected customers	Once you know who they are	Communications lead, with legal	What happened to them, what you're doing, what they should do
Regulators	Per the clocks in Section 9	Legal and privacy	The facts required, on time
Public and media	Only if needed	Communications lead	Holding statement, then facts

SECTION 9

Regulatory clocks

Several regimes may apply to the same AI incident at once, and they start their clocks from different events: awareness, detection, classification or a materiality determination. Identify every potentially applicable trigger immediately and confirm it with counsel. Initial reports can often be incomplete and updated later.

Regime	Applies to	Clock starts	Deadline
EU AI Act, Article 73 ^[1]	Providers of high-risk AI systems, for serious incidents. Deployers have related notification duties under Article 26(5) that should be assessed separately.	Awareness of the incident and a causal link, or its reasonable likelihood	No later than 15 days. 2 days for a widespread infringement or serious and irreversible disruption of critical infrastructure. 10 days where a person may have died. An initial incomplete report is allowed.
EU GDPR, Article 33 ^[2]	Personal data breaches under EU GDPR	Awareness of the breach	Without undue delay and, where feasible, within 72 hours, to the supervisory authority
UK GDPR, Article 33 ^[3]	Personal data breaches under UK GDPR	Awareness of the breach	Without undue delay and, where feasible, within 72 hours, to the ICO
NIS2, Article 23 ^[4]	Essential and important entities in the EU, for significant incidents	Awareness of the significant incident	Early warning within 24 hours, incident notification within 72 hours, final report within one month
DORA, Article 19 and Delegated Regulation (EU) 2025/301 ^[5]	EU financial entities, for major ICT-related incidents	Classification as major (and awareness)	Initial notification within 4 hours of classification and no later than 24 hours after awareness; intermediate report within 72 hours of the initial notification; final report within one month of the latest intermediate report
US SEC, Form 8-K Item 1.05 ^[6]	US-listed public companies, for material cybersecurity incidents	Determination that the incident is material (made without unreasonable delay)	Within 4 business days of that determination

Check with counsel. Scope, application dates, thresholds and national rules vary, and some EU AI Act obligations are phased in over time. Verify against the current consolidated texts linked in the notes. This table is a starting point for the conversation with your lawyers, not legal advice (see Terms of use).

SECTION 10

Restarting safely

The pressure to restore the AI to full authority will be strong. Agree these criteria before anyone does:

- ☐ The entry point is understood and closed.
- ☐ Every action in the incident window is accounted for and, where possible, undone.
- ☐ The control that failed has been fixed **and tested**, not just changed.
- ☐ Containment has been tested again.
- ☐ Two named people approve the restart, with the evidence in front of them.
- ☐ The system restarts with **less** authority than before (lower limits, narrower tools, more step-ups), and earns the rest back.

SECTION 11

Briefing the board

Boards will ask five questions. Have one page that answers them:

1. **What happened, in plain words?** One paragraph, no jargon.
2. **Is it over?** What is contained, what is still at risk.
3. **What is our exposure?** Money, data, customers, regulators, reputation.
4. **Who is accountable?** The named owner of the system and of the response.
5. **What are we changing, and by when?** Specific changes with dates and owners.

Add the timeline of decisions from your incident log. Boards trust a response more when they can see who decided what, and when.

SECTION 12

Readiness check

Answer yes or no. Be strict: "we're working on it" is a no.

#	Question	Yes / No
1	We have an inventory of AI systems, including AI features inside SaaS.	
2	Every AI system that can act has a named owner.	
3	Every AI system that can act has containment options that were tested in the last six months.	
4	We can list every action an AI system took in a given hour, from logs, without asking the AI.	
5	Irreversible or high-consequence AI actions are blocked or need a pre-agreed step-up.	
6	No AI system shares credentials with a person or holds broad admin rights.	
7	We know which AI systems read untrusted content (email, web, uploaded files).	
8	We pin model versions or re-test when the model changes.	
9	Our AI vendor contracts cover training on our data, sub-processors, logs and incident notification.	
10	We know which processes stop if a model or AI provider is unavailable, and have a tested fallback.	
11	No payment or data transfer can be approved on a call or message alone.	
12	The roles table in Section 4 is filled in with names and backups.	
13	We have rehearsed an AI incident with the leadership team in the last year.	

- **11 to 13 yes: strong baseline.** Rehearse it to prove it.
- **7 to 10 yes: partial readiness.** You may contain the incident but struggle to size, explain or recover from it.
- **0 to 6 yes: high exposure.** Start with questions 1, 3, 4 and 12.

SECTION 13

Go further: build your organisation's plan

This handbook gives you the generic version. The next step is a version built for your organisation, your systems and your people.

The **AI Crisis Readiness Review**, offered by Keystone Executive (Ankush Chowdhary's consulting practice), is an independent, fixed-scope assessment of up to three of your critical AI systems over 10 business days. It delivers:

- **Customised roles and decision rights:** who contains, who undoes, who notifies, who approves a restart.
- **Containment options for each system:** what each one supports today, and what it should.
- **Your first-hour actions:** a first-hour and first-day procedure written for your systems.
- **Notification triggers:** the regimes that may apply to you and what starts each clock, as input for your counsel.
- **Restart criteria:** how each system earns its authority back after an incident.
- **A prioritised 30-day plan:** your five most dangerous gaps and the actions that close them, with an executive readout.

Find out more at aicontrolarchitecture.org/ai-crisis-handbook, alongside the free ACA framework, requirements catalogue and templates.

Once the plan exists, a tabletop exercise tests whether the organisation can execute it.

Notes and sources

1. Regulation (EU) 2024/1689 (AI Act), Articles 73 and 26(5). The official consolidated text is the controlling source. eur-lex.europa.eu/eli/reg/2024/1689/oj
2. Regulation (EU) 2016/679 (GDPR), Article 33. eur-lex.europa.eu/eli/reg/2016/679/oj
3. UK GDPR, Article 33; ICO guidance on reporting a breach. ico.org.uk/for-organisations/report-a-breach
4. Directive (EU) 2022/2555 (NIS2), Article 23. eur-lex.europa.eu/eli/dir/2022/2555/oj
5. Regulation (EU) 2022/2554 (DORA), Article 19, and Commission Delegated Regulation (EU) 2025/301, Article 5. [eur-lex.europa.eu: Delegated Regulation 2025/301](https://eur-lex.europa.eu/eli/reg/2025/301/oj)
6. US SEC Release No. 33-11216, Form 8-K Item 1.05. [sec.gov: fact sheet](https://www.sec.gov/foia33/2025/20250301.htm)
7. Replit agent incident (July 2025), as reported by Jason Lemkin and covered by eWeek and others. eweek.com
8. Moffatt v. Air Canada, 2024 BCCRT 149. Summary by the American Bar Association. [americanbar.org](https://www.americanbar.org)
9. Arup deepfake fraud: Hong Kong Police statements (February 2024) and Arup confirmation (May 2024), reported by CNN. [CNN via abc17news.com](https://www.cnn.com)

Terms of use

This handbook is provided free of charge, **as is**, without warranty of any kind, express or implied, including fitness for a particular purpose, accuracy or completeness. By downloading or using it, you accept that the publisher and authors are not liable for any loss or damage arising from its use or from decisions made in reliance on it.

It is general information, not legal, regulatory or professional advice. Laws, deadlines and obligations change and depend on your jurisdiction, sector and circumstances. You are responsible for your own due diligence and for confirming your obligations with qualified advisers. Company names and incidents are cited from public reporting for illustration only.

The AI Crisis Preparedness Handbook, free edition, v1.2, October 2026. Written by Ankush Chowdhary, creator of the AI Control Architecture. Published by the AI Control Architecture project. Licensed under CC BY 4.0: share and adapt with attribution.