AI agents inside the risk and compliance department

Who gathers, who judges, and what is never rated, passed, closed, signed or filed on an agent's word

IN BRIEF

01 The risk and compliance department is the second line. It exists to stand apart from the work it checks, and a named person answers for its judgements. That shapes what an agent may do there more than in any other department.

02 Three things are on record. First, a control can fail silently. In a 2026 enforcement case a firm's surveillance system stopped receiving orders and raised no alerts for them. The regulator found it, not the firm.

03 Second, overstating AI draws enforcement. A securities regulator has charged firms, and taken penalties from them, for claiming AI they did not use.

04 Third, the regulators have chosen accountability over new AI rules. Senior managers answer. One country's model risk guidance now says generative and agentic AI are outside its scope.

05 No vendor page we opened says a control is marked or a report filed with no person. Several say a professional approves. The public templates are looser. Of 15 we opened, 2 hold every action for a person.

06 Hosting is not the difference. More than one vendor offers its risk software on the customer's own infrastructure.

07 This paper is for the second line and its work with internal audit: regulatory change, policies, the risk register, controls and testing, incidents and notifications, third parties, conduct cases, audit and reporting. Financial crime checks in banks, the legal department and security operations are not the subject. It ends with seven questions for any vendor and a place to start.

08 The position throughout is simple, and it has four parts. Agents gather, compare and draft, and a named officer judges and signs. An agent says what it read and what it could not read. A check that found nothing is reported with its scope. And every call is recorded. No agent rates a risk, marks a control as passed or failed, closes a finding, signs an attestation, decides whether an incident must be notified, or files or sends anything to a regulator, an auditor or a customer.

IN THIS PAPER

01 Where risk and compliance departments are today 2
02 What the public templates do 3
03 Three doors, one gate 4
04 What an agent may do 5
05 What the rules already ask 7
06 A worked example: one insurer's risk and compliance department 9
07 Seven questions before an agent touches a register 10
08 Where to start 10

PART 01

Where risk and compliance departments are today

2%

of firms in a 2026 survey by the Bank of England and the Financial Conduct Authority reported agents with broad autonomy. 80% reported agents that recommend while people decide

A control can fail and say nothing. On 24 March 2026 the United Kingdom's Financial Conduct Authority fined Dinosaur Merchant Bank Limited £338,000 [1].

The firm had brought in a new order system. The final notice says "alerts were not being generated by its automated surveillance system" for orders and trades placed through it, "as those orders and trades were not being ingested into the surveillance system for review. This had been the case since the new order system had been introduced" [1]. The gap ran from 1 June to 8 October 2024.

The notice says how it came to light. "In September 2024, the Authority identified potentially suspicious trading activity conducted by two of DMBL's clients which may have met the threshold for reporting but for which the Authority had not received any STORs from DMBL." A STOR is a suspicious transaction and order report. The firm found the cause "in response to the Authority's enquiries" [1].

"The failings impacted 2,194 trades", the notice says [1].

What that teaches. The system raised no alert because it received no orders. From the inside, that looked the same as a quiet period. So every agent in this paper is instructed to say what it read and what it could not read.

Overstating AI draws enforcement. On 18 March 2024 the United States Securities and Exchange Commission announced settled charges against two investment advisers, Delphia (USA) Inc. and Global Predictions Inc., "for making false and misleading statements about their purported use of artificial intelligence (AI)". The firms agreed to "pay $400,000 in total civil penalties" [2].

The release says the order against the first finds that it "did not in fact have the AI and machine learning capabilities that it claimed". Both settled "without admitting or denying" the findings. The Commission's chair said: "Investment advisers should not mislead the public by saying they are using an AI model when they are not" [2].

A risk department is the last place to overstate what an agent does. This paper claims four verbs for its agents: they read, compare, list and draft.

The regulators have chosen accountability over new rules. The Financial Conduct Authority's own page says: "We do not plan to introduce extra regulations for AI. Instead, we'll rely on existing frameworks" [4]. Its October 2026 paper on good practice says of senior managers that they "should, however, take reasonable steps to understand the risks arising from AI use within their area of responsibility" [5].

In the United States, the banking agencies replaced their 2011 model risk guidance on 17 April 2026. The bulletin that carries the new guidance says: "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance" [6] [7]. Part 5 has the detail.

So there is no rulebook for an agent to pass. There is a person who answers.

What firms report. The Bank of England and the Financial Conduct Authority surveyed financial firms in the United Kingdom and published the results on 26 October 2026. 196 firms responded. "86% of respondents reported using artificial intelligence (AI), up from 75% in 2024." The report warns that the two years' samples differ [3].

On agents: "80% reported having AI agents in production that provided decision support (ie, AI agents recommend actions, humans decide and execute). Only 2% reported deploying agents with broad autonomy" [3]. The survey covers one sector in one country.

What the vendors say. We opened no vendor page that says a control is marked as passed, a finding closed, a report filed or a questionnaire sent with no person. Several say a professional approves.

- Thomson Reuters, on its tax and compliance agents: "a professional approves the return before anything is filed" [8].
- Archer, on its agents for governance, risk and compliance: "a reviewer approves the result before it reaches your controls" [9].
- LogicGate describes its agents as working "with humans always in the loop as the final approvers" [10].
- DataSnipper, on its audit agents: "Your team collaborates, reviews, and signs off" [11].
- Optro says "all AI generated deliverables go through a practitioner-mediated, human review process before being accepted for use in your environment" [12].
- Drata says its questionnaire agent works "while keeping users firmly in control of all critical actions" [13].

Some pages describe steps that run with no person. Each vendor sets a limit on them.

- ServiceNow, on a product it expects to release in early 2027: "When something breaks, ServiceNow flags it, opens a ticket, and can remediate automatically within defined policies" [15].
- Vanta: "Vanta tests your controls automatically and flags gaps as they appear" [16].
- Deloitte says its controls product "Automates significant portions of the operating effectiveness testing of a control", and describes it as "Led by human intelligence and enhanced through AI" [17].
- MetricStream's home page shows a demonstration in which some steps run without a person and one waits: "Tier 1 requires human approval" [14].

These are design choices, and each vendor explains them. A buyer should know which steps run alone.

PART 01 · CONTINUED

Hosting is not the difference. Archer says its product can be deployed “standalone or alongside Archer GRC, on-premises or SaaS” [9]. IBM says OpenPages is “Available on any cloud or on-premises” [18]. Neither page says where the agent features run. Where an agent runs is a question to ask, and more than one vendor can answer it.

PART 02

What the public templates do

2 of 15

public risk and compliance templates and examples we opened hold every action for a person

We opened 15 public templates and examples for risk, compliance and audit work. They come from two workflow platforms, one cloud provider’s pattern library, one software vendor’s guidance and gallery, one trade article and one agent framework’s documentation. We read each for one thing: does a person approve before a register is written, a ticket opened, a status set or something sent out?

WHAT THE TEMPLATE DOES	DOES A PERSON APPROVE FIRST?
Reads a regulator’s publications, rates the urgency, then “Creates a high-priority or normal-priority compliance ticket in your ticketing system”	**No step**
Scores a supplier from public pages, then “appends or updates the vendor entry” in a risk register	**No step**
Rates new legislation. “auto-approved items are logged directly”	**Only a notice for the flagged ones**
Drafts answers to a security questionnaire. Each draft “passes an email approval gate with approve/reject links”	**In part. The page does not say who sends the answers**
Pauses a case and waits. “The graph literally cannot proceed past the interrupt node without external input”	**Yes. This one is a code pattern in a trade article**

The count. 2 of the 15 hold every action for a person. Both are code patterns, not ready templates, and one is an example for financial crime alerts. 5 do so in part: for some items, for some actions, or as a stated practice. 8 have no step at all [19] [20].

Two findings matter more than the count. First, the ready-made templates write before anyone looks. They rate a supplier and update the register. The person is told afterwards.

Second, the rating itself is the agent’s. In several templates the model sets the urgency, the tier or the status, and the workflow routes on it. In one, the word “approval” names a status the workflow assigns, not a person’s act [19]. That is the pattern this paper avoids: no agent here rates anything.

One design stands out. A guidance page for governing a company’s own agents lists its options plainly. One is “No automated enforcement”: an administrator is told and acts. Another deletes, and the page says “This action can’t be undone” [20]. The choice is visible. Ask every product to show the same.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a risk and compliance department that means three kinds of thing.

Documents.

Policies and procedures, the control library, the obligation register, past audit reports, regulatory correspondence and committee papers. These are indexed once as knowledge sources, with each reader's permissions kept, so a colleague does not reach through an agent a case or a finding they could not open themselves.

Systems of record.

The risk and control system, the incident log, the case system, the training system. Most are reached through an interface the organisation already runs. The organisation gives the platform the description of that interface and picks which operations become tools.

Files and outside sources.

A new rule, a control test file, an incident report or a supplier's questionnaire arrives as a file. Document stores, tickets, reporting databases and web search are reached through MCP servers deployed in the organisation's own cluster.

All three pass one gate: every tool call is checked against the organisation's policy, which can allow it, record it, hold it for a person's approval, or block it.

FOUR WAYS INTO A RISK AND COMPLIANCE DEPARTMENT'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A new rule, a control test file, an incident report, a supplier's questionnaire, an export of the risk register	Nothing is connected. The safest start
02	**A knowledge source**	Policies, the control library, the obligation register, past audit reports, committee papers	Permissions are kept per reader
03	**An imported description of an interface**	The risk and control system, the incident log, the case system, the training system	The organisation picks the operations. Changes wait for approval
04	**An MCP server, deployed in the organisation's cluster**	Document stores, tickets, reporting databases, web search	It runs inside. Check what each tool does

PART 03 · CONTINUED

8,667

operations sent as changes, out of 16,041, in the public interface descriptions we imported (111 files read)

What is published. We ran Katonic's importer on 111 public interface descriptions from 56 publishers of systems and public sources a risk and compliance department uses. 103 import as they are and 8 are over the 5 MB import limit. None was refused. The 103 give 16,041 operations. 7,374 read, and 8,667 are sent as changes. The method alone cannot tell which of those only read. The counts are per file, and some publishers split one system into several files or list several versions [21].

Who the publishers are. Identity governance and access review tools, e-signature and attestation, document and records stores, reporting tools, a learning platform, ticket and wiki systems, one compliance automation vendor, and public sources: legislatures, a deposit insurer, company registries, a sanctions data service and a vulnerability programme.

What that exercise taught us. More than half of the operations change something. That is the reason for the approval rule, and for giving an agent the read operations and nothing else. So pick the operations at import, read what each one does, and start with the ones that only read. A description over the limit is imported in the part the organisation needs.

The systems at the centre publish nothing. The governance, risk and compliance suites publish no description a file address reaches. Neither do most regulators' registers and legislation services. For those, the description comes from the organisation's own installation, or the agent reads an export or an uploaded file. Our search for the suites was light, and a miss does not prove that no description exists behind a sign-in.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named risk or compliance officer decides.

01	**It reads** the rules, policies, registers, evidence and records it has been given, through read tools only.
02	**It compares** a policy with a rule, or evidence with a control description, and gives the paragraph on each side.
03	**It lists** what is missing, late, unowned or at odds.
04	**It drafts** the impact note, the log entry, the working paper, the committee report. A draft is a file for the officer. It is not a record.
05	**It says what it read and what it could not read,** and it says when it found nothing, with the scope of the check. It is instructed not to fill a gap from memory.
06	**It stops** before anything that changes a register, a case or a filing. A tool that changes something waits for approval unless the organisation has decided otherwise for that tool.

PART 04 · CONTINUED

That gives 35 agents in nine areas. The full list is a separate document.

AREA	AGENTS	WHAT THEY PREPARE	WHO DECIDES
Regulatory change	4	Which new publications bear on the organisation; the policies and controls a rule touches; draft register entries; a draft consultation response	The compliance officer and the head of compliance
Policies	4	Answers with the paragraph cited; requirements a policy does not address; a draft policy change; who has not attested	The policy owner and the approving committee
Risk register and assessments	4	Register entries with gaps; an assessment pack; the indicator report; a risk event summary	The risk owner and the risk officer
Controls and testing	4	Evidence set beside each test step; what the evidence does not show; a proposed control mapping; overdue findings	The tester, the control owner, the audit team
Incidents and notifications	4	A draft log entry; the clocks the procedure gives; a draft notification; the breach register summary	The incident owner and the accountable officer
Third parties	3	An assessment pack; register entries with gaps; draft questionnaire answers	The risk officer and the responsible officer
Conduct and cases	4	A case intake summary; an investigation file summary; late declarations; overdue training	The case handler, the investigator, the compliance officer
Audit	4	A planning pack; an evidence request list; a draft working paper; a draft finding	The audit lead and the auditor
Reporting	4	A regulatory return for review; a draft committee report; the status of the compliance plan; reports and memos	The accountable officer and the chief risk or compliance officer

Four lines are drawn tighter than the rest.

The judgement.

An agent may set the evidence beside each step of a test and list the steps with no evidence. It gives no pass and no fail. It proposes no risk rating. The compliance or audit team closes a finding. An agent that marked its own work as passed would not be a second line.

The regulator.

An agent may draft a notification or a return in the regulator's form and mark every field the record does not fill. The accountable officer approves and files it. For an incident, the Notification Deadline Check lists the clocks the organisation's own procedure gives. It does not decide whether the incident must be notified, and it notifies nobody.

The reporter.

A case agent reads what the case handler is entitled to read, and no more. No agent tries to identify the person who made a report. An investigation summary draws no conclusion about a person. The lists of overdue training and missing attestations are lists of records. They score nobody.

The silence.

Every agent is instructed to report what it could not open, and to give the scope of a check that found nothing. Every call is recorded, with who asked, which tool and what was decided.

PART 05

What the rules already ask

This part needs a read by the organisation's own lawyers and auditors before it is relied on. It reports what the published texts say; it is a summary, not legal advice. Nothing in it says that a platform makes an organisation compliant. Whether a rule applies depends on the organisation's sector, size and country.

Three lines, and evidence a person can stand behind

The three lines model. The Institute of Internal Auditors' model puts independence with internal audit: "Internal audit's independence from the responsibilities of management is critical to its objectivity, authority, and credibility." The second line sits with management: "Management's responsibility to achieve organizational objectives comprises both first and second line roles." In the model, "Second line roles assist with managing risk" [22].

Internal audit standards. "Internal auditors must document information and evidence to support the engagement results." The work must be documented so that a competent person "could repeat the work and derive the same engagement results" [23].

External audit standards in the United States. Audit documentation "must contain sufficient information to enable an experienced auditor, having no previous connection with the engagement" to understand the work and to determine who performed it [24].

What follows. None of these texts, in the parts we read, mentions AI. They ask whether the work can be repeated and who did it. An agent's draft meets them only through the person who checks it and signs.

A named person answers

The United Kingdom. The law lets a regulator act against a senior manager who "did not take such steps as a person in the senior manager's position could reasonably be expected to take to avoid the contravention occurring (or continuing)". The regime aims at "making individuals more accountable for their conduct and competence" [25].

The same regulator on AI. It plans no extra rules [4]. Its good practice paper says: "Effective oversight is not simply about reviewing AI outputs, but about ensuring people have the right authority, information, capability and escalation routes to challenge or intervene where needed" [5]. The paper is a set of examples, not rules.

Notification clocks

Personal data. A controller notifies the supervisory authority "without undue delay and, where feasible, not later than 72 hours after having become aware of it" [26].

Financial firms in the European Union. For a major incident, the first report is due "within four hours from the classification of the ICT-related incident as a major ICT-related incident and no later than 24 hours from the moment the financial entity has become aware" of it [27].

Network security. An early warning is due "within 24 hours of becoming aware of the significant incident", a notification "within 72 hours", and "a final report not later than one month" after that [28]. The duty reaches an organisation through national law.

Why no agent decides. Whether an event counts, and when the organisation became aware of it, are judgements. The agent lists the clocks from the organisation's own procedure. The officer decides.

Model risk guidance and its scope

The United Kingdom. The prudential regulator defines the term: "Model risk is the potential for adverse consequences from model errors or the inappropriate use of modelled outputs to inform business decisions." Its policy took effect on 17 May 2024, and it makes senior managers accountable for remediation [29]. The text we searched does not mention AI by name.

The United States. A supervisory letter of 17 April 2026 issued revised guidance "which supersedes and replaces SR letter 11-7", the 2011 letter [6]. The banking regulator's bulletin of the same date says generative and agentic AI models "are not within the scope of this guidance". It adds that the guidance "does not set forth enforceable standards or prescriptive requirements", and that the agencies plan a request for information [7]. We did not read the guidance itself, only the letter and the bulletin.

PART 05 · CONTINUED

What one prosecutor's guide asks

The United States Department of Justice's guide for prosecutors, updated in September 2024, asks about a company's own use of AI. "What baseline of human decision-making is used to assess AI? How is accountability over use of AI monitored and enforced?" It asks about AI "in its commercial business and in its compliance program" [30]. A compliance programme that uses an agent is inside the question.

A whistleblower's identity

The European Union's directive says: "Member States shall ensure that the identity of the reporting person is not disclosed to anyone beyond the authorised staff members competent to receive or follow up on reports, without the explicit consent of that person" [31]. That is why a case agent reads only what the handler may read, and why none tries to identify a reporter.

The AI Act

There is no entry called compliance, risk or audit. One entry can still matter. The Act's list covers systems meant "to monitor and evaluate the performance and behaviour of persons" at work. Where a system is high-risk, "Deployers shall assign human oversight to natural persons who have the necessary competence, training and authority" [32]. Whether an entry catches a particular use is for counsel.

The dates have moved. Regulation (EU) 2026/1744 says the high-risk rules apply from "2 December 2027 as regards AI systems classified as high-risk pursuant to Article 6(2) and Annex III" and from 2 August 2028 for those under Article 6(1) and Annex I [33].

PLACE	WHAT WE READ	WHAT IT ASKS
Internal audit, worldwide	Global Internal Audit Standards	Work that a competent person "could repeat" [23]
United Kingdom, financial firms	The duty of responsibility	Reasonable steps by a named senior manager [25]
European Union	Data protection, operational resilience and network security law	First notices in 72, 24 and 24 hours [26] [27] [28]
United States, banks	The 2026 model risk letter and bulletin	Generative and agentic AI "not within the scope" [7]

We did not open the international auditing standards, the internal control framework most audits rely on, the texts of the ISO standards for risk, compliance and AI management, the accountability regimes of Singapore, Hong Kong and Australia, or any Gulf regulator's guidance on AI. An organisation should check each with its own advisers.

PART 05 · CONTINUED

What the rules have in common

01	**A person answers.** The texts name senior managers, management bodies, auditors and authorised staff.
02	**Evidence must be repeatable and attributable.** Someone else must be able to see what was done and who did it.
03	**The clock and the judgement belong to the organisation.** No text we read lets a tool decide what is notified.
04	**Claims about AI must be true.** That includes claims made inside the organisation.

PART 06

A worked example: one insurer's risk and compliance department

Take a made-up insurer: 2,600 staff, a risk and compliance department of 45, and an internal audit team of 9. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

On a new rule, the regulator publishes a paper of 84 pages. The agent lists 7 policies and 12 controls that refer to the same subject, with the passage in each, and 3 requirements that none of them covers. The compliance officer decides the impact and assigns the changes.

02

Before a control test, the tester opens a test with 9 steps. The agent sets the evidence beside each step and lists 2 steps with no evidence. It also reports that it could not open 1 of the 14 files. The tester obtains the file, performs the test and records the result.

03

After an incident, a supplier reports a data loss at 16:40. The agent drafts the log entry and lists the three notifications the insurer's own procedure names, to whom, and the time each clock would end. The compliance officer decides which apply. The accountable officer approves and files.

04

Before the committee, the agent drafts the quarterly report from the registers. One check found nothing: no overdue findings in the two registers it read, for the quarter. It says so, and names the third register it has no access to. The chief risk officer asks for that one and approves the report.

Every one of these changes nothing by itself, rates nothing, files nothing, and leaves a record of who asked, which tool ran and what was decided.

PART 07

Seven questions before an agent touches a register

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the vendor's service? Where is the model, and do our case files and regulatory correspondence leave? More than one vendor can run on your side. Ask the next six of each.

02 **What can it read?**
Does it keep each reader's permissions, or does it read every case and finding as one all-seeing account?

03 **What can it change?**
List every tool that rates, marks, closes, attests, files or sends. Is each one separate from the read tools, and does each wait for a person?

04 **Who answers for it?**
Which named officer owns each agent, and what did they check before it was used?

05 **What does it say when it finds nothing?**
Does the output give the sources and the period it covered, and what it could not open?

06 **What is recorded?**
Who asked, which tool ran, what was decided, for how long is that kept, and can an auditor read it?

07 **Can we stop it?**
One agent, one tool or one server, at once, by us. And is everything we say about it, to the board, a customer or a regulator, what it does?

PART 08

Where to start

Pick work where reading is slow, the judgement stays with an officer, and nothing reaches a regulator on the agent's word.

Questions and answers on the policies
Colleagues ask what the policy says and get the answer with the paragraph cited. No register is touched.

Regulatory change impact note
The agent lists the policies and controls a new rule touches. The compliance officer decides the impact.

Control test preparation
The agent sets the evidence beside each step of the test. The tester tests.

One team, one rule or one control, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/RISK-COMPLIANCE

SOURCES — Read on 23 November 2026. Figures are the publishers' own and were not audited by us.

01 Financial Conduct Authority, Final Notice to Dinosaur Merchant Bank Limited, 24 March 2026. https://www.fca.org.uk/publication/final-notices/dinosaur-merchant-bank-limited-2026.pdf

02 United States Securities and Exchange Commission, press release 2024-36, "SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence", 18 March 2024. https://www.sec.gov/newsroom/press-releases/2024-36

03 Bank of England and Financial Conduct Authority, "Artificial intelligence in UK financial services - 2026", published 26 October 2026. https://www.bankofengland.co.uk/report/2026/artificial-intelligence-in-uk-financial-services-2026

04 Financial Conduct Authority, "AI and the FCA: our approach", first published 8 September 2025, last updated 18 November 2026. https://www.fca.org.uk/firms/innovation/ai-approach

05 Financial Conduct Authority, "AI in financial services: good practice and areas for improvement", published 19 October 2026, last updated 20 October 2026. https://www.fca.org.uk/publications/good-poor-practice/ai-financial-services

06 Board of Governors of the Federal Reserve System, SR 26-2, "Revised Guidance on Model Risk Management", 17 April 2026. https://www.federalreserve.gov/supervisionreg/srletters/sr2602.htm

07 Office of the Comptroller of the Currency, OCC Bulletin 2026-13, "Model Risk Management: Revised Guidance", 17 April 2026. https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html

08 Thomson Reuters, press release of 17 November 2026 on agents in its tax and compliance products. https://www.thomsonreuters.com/en/press-releases/2026/november/thomson-reuters-brings-cocounsel-agentic-ai-to-the-tax-and-compliance-products-professionals-trust

09 Archer, home page, undated, read 23 November 2026. https://www.archerirm.com/

10 LogicGate, home page, undated. https://www.logicgate.ai/

11 DataSnipper, home page, undated. https://www.datasnipper.com/

12 Optro, page on AI in its platform (frequently asked questions), undated. https://optro.ai/platform/ai/

13 Drata, press release of 24 March 2026. https://drata.com/about/news/drata-unveils-industry-first-agentic-ai-capabilities-to-transform-enterprise-trust-workflows

14 MetricStream, home page (the demonstration panel), undated. https://www.metricstream.com/

15 ServiceNow, press release of 10 November 2026; general availability expected in the first quarter of 2027. https://newsroom.servicenow.com/press-releases/details/2026/ServiceNow-Brings-Agentic-Monitoring-to-Every-Control-in-the-Enterprise-at-Machine-Speed/

16 Vanta, product page, undated. https://www.vanta.com/products/agentic-trust-controls

17 Deloitte, press release of 5 August 2026 on ControlCatalyst.AI. https://www.deloitte.com/us/en/about/press-room/deloitte-announces-launch-of-controlcatalyst-ai.html

18 IBM, OpenPages product page, undated. https://www.ibm.com/products/openpages

19 Eight workflow templates on n8n: 13157 (policy checks and approvals), 17947 (changes on regulators' pages), 19133 (regulatory publications to tickets), 18503 (new supplier risk records), 19118 (supplier risk register), 13451 (legislation routing), 14436 (policy review tracks) and 21931 (security questionnaire answers). https://n8n.io/workflows/13157 ; https://n8n.io/workflows/17947 ; https://n8n.io/workflows/19133 ; https://n8n.io/workflows/18503 ; https://n8n.io/workflows/19118 ; https://n8n.io/workflows/13451 ; https://n8n.io/workflows/14436 ; https://n8n.io/workflows/21931

SOURCES continued

20 Seven further templates and examples: Zapier, contract compliance review template. https://zapier.com/templates/automated-contract-compliance-review-risk-assessment ; Amazon Web Services Prescriptive Guidance, a pattern for configuration compliance with an agent (read in the German machine translation; the English address did not load). https://docs.aws.amazon.com/de_de/prescriptive-guidance/latest/patterns/streamline-amazon-ec2-compliance-management-with-amazon-bedrock-agents-and-aws-config.html ; Microsoft Learn, Compliance Hub guidance for Copilot Studio. https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/kit-compliance-hub ; Microsoft, Copilot Studio gallery sample. https://microsoft.github.io/copilot-studio-gallery/solutions/agent-dlp-security-guard/ ; Microsoft Learn, an internal audit case study. https://learn.microsoft.com/power-platform/guidance/case-studies/grupo-bimbo-global-audit ; C# Corner, a trade article with a code pattern, 19 June 2026. https://www.c-sharpcorner.com/article/combining-langgraph-and-rag-for-enterprise-compliance/ ; CrewAI documentation, human input page. https://docs.crewai.com/en/learn/human-in-the-loop

21 111 public interface descriptions from 56 publishers, run through the platform's importer on 23 November 2026 (Katonic research); and the search for publishers with no downloadable description (Katonic research).

22 The Institute of Internal Auditors, "The IIA's Three Lines Model", 2020 (the file carries "Updated September 2024"), Principles 3 and 5. https://www.theiia.org/globalassets/documents/resources/the-iias-three-lines-model-an-update-of-the-three-lines-of-defense-july-2020/three-lines-model-updated-english.pdf

23 The Institute of Internal Auditors, Global Internal Audit Standards, published 9 January 2024, Standard 14.6. https://www.theiia.org/globalassets/site/standards/editable-versions/globalinternalauditstandards_2024january9_editable.pdf

24 Public Company Accounting Oversight Board, AS 1215 "Audit Documentation", paragraph .06A, and AS 1105 "Audit Evidence". https://pcaobus.org/oversight/standards/auditing-standards/details/AS1215 ; https://pcaobus.org/oversight/standards/auditing-standards/details/AS1105

25 Financial Services and Markets Act 2000, section 66A(5). https://www.legislation.gov.uk/ukpga/2000/8/section/66A ; and Financial Conduct Authority, "Senior Managers and Certification Regime", last updated 24 April 2026. https://www.fca.org.uk/firms/senior-managers-certification-regime

26 Regulation (EU) 2016/679 (General Data Protection Regulation), Article 33(1). http://publications.europa.eu/resource/celex/32016R0679

27 Regulation (EU) 2022/2554 (Digital Operational Resilience Act), Article 19; and Commission Delegated Regulation (EU) 2025/301 on the time limits for incident reports. http://publications.europa.eu/resource/celex/32022R2554 ; http://publications.europa.eu/resource/celex/32025R0301

28 Directive (EU) 2022/2555 (NIS2), Article 23. http://publications.europa.eu/resource/celex/32022L2555

29 Prudential Regulation Authority, Supervisory Statement SS1/23, "Model risk management principles for banks", in effect 17 May 2024. https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2023/ss123.pdf

30 United States Department of Justice, Criminal Division, "Evaluation of Corporate Compliance Programs (Updated September 2024)". https://www.justice.gov/criminal/criminal-fraud/page/file/937501/dl

31 Directive (EU) 2019/1937 on the protection of persons who report breaches of Union law, Article 16(1). http://publications.europa.eu/resource/celex/32019L1937

32 Regulation (EU) 2024/1689 (AI Act), Article 26(2) and Annex III point 4(b). http://publications.europa.eu/resource/celex/32024R1689

33 Regulation (EU) 2026/1744 of 8 July 2026, the replacement of Article 113(c) of Regulation (EU) 2024/1689. http://publications.europa.eu/resource/celex/32026R1744