AI agents inside the operations department

Who prepares, who decides, and what never reaches a person outside on an agent's word

IN BRIEF

01 Operations is the department where automation is oldest. Case processing, forms, letters and records were run by software long before anyone spoke of agents.

02 It is also where the best documented failures are. In Australia, a debt scheme averaged people's income and raised debts with, in the Royal Commission's words, the aim "to reduce (and preferably obviate) human involvement". In the United Kingdom, hundreds of people who ran post offices were, the public inquiry says, "wrongly accused of theft and fraud due to faults in the Horizon IT system". In the Netherlands, a parliamentary committee found that the ministry "administered childcare allowance as a mass process".

03 None of these was AI. Each was plain automation, or plain software, run at scale with nobody placed to stop it. They show what an agent must not be allowed to become.

04 Most vendors of automation place a person at the exceptions. One vendor's own guidance tells builders: "Don't rely on instructions" to "prevent unauthorized actions". The public templates are less careful. Of 15 we opened, 2 hold every item for a person.

05 Hosting is not the difference. Several automation vendors offer agents on the customer's own infrastructure, and this paper quotes them.

06 This paper is about the chief operating officer's organisation and its shared services: case processing, documents, procedures, quality, exceptions, facilities, continuity, service providers and reporting. It ends with seven questions for any vendor and a place to start.

07 The position throughout is simple, and it has five parts. Agents prepare and a named case officer decides. Read tools are kept apart from change tools. Change tools are held for approval. Every call is recorded. And nothing reaches a person outside on an agent's word. No agent decides a case, approves, rejects, closes a case, scores or ranks a case or a person, recommends an outcome, changes the record of a decision, or sends a decision to anyone.

IN THIS PAPER

01 Where operations 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 operations department 9
07 Seven questions before an agent touches a case 10
08 Where to start 10

PART 01

Where operations departments are today

433,000

people against whom one government "unlawfully asserted" debts worked out by income averaging, as a court recorded in 2021. Plain automation, not AI

The lesson of this department is older than AI. The three best documented failures in back-office work were not AI. We report each in the words of the body that examined it.

Australia: debts raised by averaging. The Royal Commission into the Robodebt Scheme reported on 7 July 2023. The scheme, it says, "relied heavily on a process known as" income averaging "to assess income and entitlement to benefit". The Commission's finding: "As used, it neither produced accurate results nor complied with the income calculation provisions of the Social Security Act 1991 (Cth)."

Where officers had once examined the file, the report says, "it was now anticipated that in most instances the entire process would be conducted online, the aim being to reduce (and preferably obviate) human involvement" at the department's end. The Commission's own summary: "It was a costly failure of public administration, in both human and economic terms" [1].

Two years earlier, the Federal Court had approved the settlement of a class action. The judge recorded that the government "unlawfully asserted such debts", "totalling at least $1.763 billion against approximately 433,000 Australians" [2].

United Kingdom: faults in a branch accounting system. The Post Office Horizon IT Inquiry's own summary of its first volume says: "hundreds of Post Office postmasters, managers and assistants were wrongly accused of theft and fraud due to faults in the Horizon IT system". The inquiry sat for "226 days of hearings" [3].

The Netherlands: an allowance run as a mass process. A committee of the Dutch House of Representatives reported in December 2020 under the title "Unprecedented injustice". It "notes that basic principles of the rule of law were breached in the administration of childcare allowance". The ministry, it says, "administered childcare allowance as a mass process" [4].

A year later the Dutch data protection authority fined the tax administration 2.75 million euros. Among its findings: the administration used applicants' nationality "as an indicator in a system that automatically designated certain applications as risky" [5].

What the three have in common. A rule, or a system's output, was applied to many people at once. The output was treated as the answer. And nobody was placed to look at one file and stop it. An agent that scores a case, or closes one, would be the same design.

Hosting is not the difference. Running agents on the organisation's own side is offered by several automation vendors, in their own words.

- UiPath says its agents and orchestration "are now available on-premises, on the Kubernetes infrastructure customers already operate on" [6].
- Automation Anywhere describes a stack run entirely on the customer's premises: "No external API calls. No data leaving the building" [7].
- Kore.ai's documentation offers a self-hosted option "For organizations requiring on-premises deployment" [8].
- IBM's product page says its agent product "Supports multiple clouds (IBM Cloud, AWS and more) or on-premise deployment" [9].

So where an agent runs is a question to ask, and more than one vendor can answer it well. The questions that separate them come after.

Most vendors place a person at the exceptions. This is the clearest pattern in the pages we opened.

- UiPath's platform page lists "Human judgment as a first-class workflow step" and says "People intervene where judgment matters" [10].
- Camunda's page asks: "Where does accountability require a person? Keep human judgment in the process" [11].
- Kore.ai's page says: "Agents act only within the limits you set. Everything above routes to a person." Reviewers, it says, see "the source documents, the checks that ran, and the agent's recommendation" [8].
- Automation Anywhere describes a business that can "operate, adapt, and optimize workflows on its own, with minimal human intervention". The same page asks readers to "Ensure human judgment remains integral to decision-making, especially in complex, high-stakes, or unforeseen scenarios" [12].
- Workato's page says: "Human decision points ensure policy compliance for discounts, exceptions, and high-stakes actions" [13].

We opened no vendor page that says a person decides every case. That suits many kinds of work. It is not what this paper recommends for a case about a person.

One vendor tells builders not to trust instructions. Microsoft's guidance for builders of agents that act without a prompt, dated 5 November 2026, says: "For high-stakes tasks, keep a human in the loop." It goes on: "Enforce strict limits on the agent's actions. Don't rely on instructions", such as a line telling the agent to send an email only after a check, "to prevent unauthorized actions" [14].

A line in a prompt is not a control. A tool the agent does not hold, or a tool that waits for a person, is.

What the published accounts leave out. We found no public account, with a primary report, of an AI agent that wrongly approved or closed a case in a back office, and we do not claim one. We opened no original survey of how many operations teams use agents, and print no survey number.

PART 02

What the public templates do

2 of 15

public operations templates and examples we opened hold every item for a person before something is changed or sent

We opened 15 public templates and examples for operations work, from three workflow platforms, one cloud provider's guidance, one vendor's template library and two developer frameworks. We read each for one thing: is there a step where a person approves before a record is changed or something is sent, and does it cover every item?

WHAT THE TEMPLATE DOES	DOES A PERSON APPROVE FIRST?
Reads invoices, checks each against a policy, and "auto-approves or raises an exception"	**Only for the exceptions. The rest post to the accounting system**
Reads invoice files and checks them against purchase orders	**Only for mismatches. "Clean invoices move straight to approval"**
Sorts chat messages into incidents and requests, and creates the tickets	**No step. The page advises sending low-confidence items to "a human review queue instead of auto-creating a ticket"**
Drafts a review of an incident, logs it and emails the draft to the person who reported it	**No**
Reads a loan application and "generates a preapproval letter"	**No step in the published design**
Investigates an IT incident and proposes a fix	**Yes. It "posts an approval request" and "waits for a reviewer decision"**

The count. 2 of the 15 hold every item for a person, and one of those two is a developer framework's guide, not an operations template. 5 have a person on some paths only. 7 have none. 1 we could not judge from the page [15] [16].

Two findings matter more than the count. First, the templates that handle invoices approve the clean ones by themselves. A person sees only what the workflow sends. Whoever wrote the test has decided every item that passes it.

Second, the test is often the model's own confidence. One template flags an invoice for review "if validation fails or confidence is low". Another suggests a figure, 0.7, below which a person should look [15].

A confidence score is a setting. Whoever sets it has written a delegation of authority.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In an operations department that means three kinds of thing.

Documents.

Procedures, work instructions, the quality manual, past audit reports, continuity plans and closed cases. These are indexed once as knowledge sources, with each reader's permissions kept, so a colleague does not reach through an agent a case they could not open themselves.

Systems of record.

The case system, the workflow engine, the document store, the facilities 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 case file, a completed form, a provider's monthly report or an inspection report arrives as a file. Document stores, tickets, reporting databases and chat channels 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 AN OPERATIONS DEPARTMENT'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A case file, a form, a provider's report, an inspection report, a backlog export	Nothing is connected. The safest start
02	**A knowledge source**	Procedures, the quality manual, audit reports, continuity plans, closed cases	Permissions are kept per reader
03	**An imported description of an interface**	The case system, the workflow engine, the document store, the facilities 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, chat channels	It runs inside. Check what each tool does

PART 03 · CONTINUED

9,196

operations sent as changes, out of 16,358, in the public interface descriptions we imported (150 files read)

What is published. We ran Katonic's importer on 150 public interface descriptions from 59 publishers of systems an operations department uses. 146 import as they are and 4 are over the 5 MB import limit. None was refused. The 146 give 16,358 operations. 7,162 read, and 9,196 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 [17].

Who the publishers are. Work and project management, workflow engines and case management, document stores and e-forms, e-signature, alerting, reporting, scheduling, wikis, surveys, document reading services and messaging.

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.

The systems closest to the department publish nothing. Quality, audit and health and safety systems, facilities and maintenance systems, emergency notification, and the large process automation suites publish no description a file address reaches. For those, the description comes from the organisation's own system, or the agent reads a report or an uploaded file. A description over the limit is imported in the part the organisation needs.

So pick the operations at import, read what each one does, and start with the ones that only read.

PART 04

What an agent may do

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

01	**It reads** the case files, forms, procedures and records it has been given, through read tools only.
02	**It drafts** the summary, the checklist result, the letter, the report, with the document and the page behind each statement.
03	**It lists** what is missing from a file, what departs from the procedure, and what falls due.
04	**It says when it found nothing.** It is built not to fill a gap from memory.
05	**It stops** before anything that changes a case, a record or a system, and before anything that leaves the building. A tool that changes something waits for approval unless the organisation has decided otherwise for that tool.
06	**It decides nothing about a person or a company.**

PART 04 · CONTINUED

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

AREA	AGENTS	WHAT THEY PREPARE	WHO DECIDES
Case intake and processing	4	The kind of case and what is missing; a summary of the file; the letter for a decision the officer has recorded	The case officer
Documents and forms	4	A proposed type for each document; fields at odds with the documents; records due for disposal	A clerk, the document owner, the records manager
Procedures and knowledge	4	Answers from the procedures; a draft procedure; what a change touches; a training note	The process owner and the team lead
Quality and process audits	4	Points where a sampled file does not show a step; audit evidence; a draft nonconformity report	The quality reviewer and the auditor
Exceptions and complaints	4	The exception queue by reason; repeat exceptions; a complaint file summary; an escalation brief	The team lead, the complaints officer, the manager
Facilities and workplace	3	The kind of request; a work order made ready; findings from an inspection report	The facilities coordinator and the maintenance planner
Continuity and incidents	4	What a continuity plan has out of date; an incident timeline; a draft notice; an exercise report	The plan owner and the incident manager
Service providers	4	Where a provider's report departs from the contract; gaps in the register; the review pack; gaps in an exit plan	The service manager and the service owner
Reporting	4	The operations report; the backlog by age; a project status draft	The head of operations and the team lead

Four lines are drawn tighter than the rest.

The decision.

An agent may summarise a case file for the officer who will decide it: who, what was asked, the facts on file with the page of each, and the open points. It gives no score, ranks no case and recommends no outcome. The officer reads the file and decides.

The quality review.

An agent may compare a sampled case with the checklist and list each point where the file does not show the step was done. The reviewer judges each point. The agent scores no case and no member of staff.

The letter.

Nothing reaches a person outside on an agent's word. An agent may draft the letter that carries a decision the officer has already recorded, with the reasons the officer gave. It adds no reason of its own. The officer approves and sends.

The record.

An agent may list records that reach the end of their retention period, and those under a hold. It deletes nothing, and it does not change the record of a decision. The records manager approves each disposal.

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, and on whether it is a public body.

A decision about a person

The data protection rule. The European Union's General Data Protection Regulation says: "The data subject shall have the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her." Where an exception applies, the person keeps "at least the right to obtain human intervention on the part of the controller, to express his or her point of view and to contest the decision" [18].

A person who only clicks is not enough. The European regulators' guidelines say: "The controller cannot avoid the Article 22 provisions by fabricating human involvement." Review, they say, "should be carried out by someone who has the authority and competence to change the decision" [19].

What the court said a "decision" is. In December 2023 the Court of Justice ruled on a credit score. It held that producing the score was itself "automated individual decision-making" where a third party "draws strongly on that probability value to establish, implement or terminate a contractual relationship with that person". The Court spoke of "the broad scope of the concept of 'decision'" [20].

That case was about credit. Its reasoning is why no agent in this paper scores, ranks or recommends an outcome on a case. If the officer nearly always follows the score, the score may be the decision. Whether it is so in a given process is for counsel.

The AI Act

There is no entry called operations. An operations use is high-risk if it falls into an entry on the Act's list. One covers systems used by or for public authorities "to evaluate the eligibility of natural persons for essential public assistance benefits and services". Another covers systems meant "to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons" at work. A third covers systems meant "to evaluate and classify emergency calls" [21]. Whether an entry catches a particular use is for counsel.

Human oversight. Those who deploy a high-risk system "shall assign human oversight to natural persons who have the necessary competence, training and authority, as well as the necessary support" [21].

The right to an explanation. A person affected by a decision taken on the output of a listed high-risk system has "the right to obtain from the deployer clear and meaningful explanations of the role of the AI system in the decision-making procedure and the main elements of the decision taken" [21].

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 [22].

Public bodies and automated decisions

Canada. The federal Directive on Automated Decision-Making requires "notice through all service delivery channels in use that the decision will be made or assisted by an automated decision system". It requires "a meaningful explanation to clients of how and why the decision was made", and a system that "allows for human involvement" [23].

The United Kingdom. The government describes its Algorithmic Transparency Recording Standard as "a standardised way for public sector organisations to publish information about how and why they are using algorithmic tools". The page says it "is mandatory for all government departments" [24].

Australia. The Royal Commission recommended that where automated decision-making is used, "there should be a clear path for those affected by decisions to seek review" [1].

PART 05 · CONTINUED

Operational resilience and outsourcing

For financial firms in the European Union, the Digital Operational Resilience Act says firms "remain fully responsible for compliance with, and the discharge of, all obligations under this Regulation and applicable financial services law". They must keep "a register of information in relation to all contractual arrangements on the use of ICT services provided by ICT third-party service providers". And: "For ICT services supporting critical or important functions, financial entities shall put in place exit strategies" [25].

The banking supervisors say the same of outsourcing. The European Banking Authority's guidelines say an institution's "management body remains responsible for that institution and all of its activities, at all times" [26]. The United Kingdom's conduct regulator says firms in scope "had until 31 March 2025 to ensure they could operate their important business services within their impact tolerances" [27].

None of these texts, in the parts we read, mentions AI agents. A regulated firm should still expect the supplier of an agent to be treated as a service provider.

Quality and continuity standards

We read the public summaries only. The standards themselves are sold, and we name no clause.

The standards body's page for the quality management standard says: "ISO 9001:2015 is no longer the current edition of the standard." It lists a 2026 edition, published in September 2026. The public summary says the standard "addresses issues such as resources, competence, awareness, communication and documented information" [28]. The business continuity standard, ISO 22301, is described as a framework to "protect against, reduce the likelihood of, and ensure recovery from disruptive incidents" [28].

PLACE	WHAT WE READ	WHAT IT ASKS
European Union	General Data Protection Regulation, Article 22	No decision "based solely on automated processing"; "human intervention" [18]
European Union	Court of Justice, case C-634/21	A score a decision "draws strongly on" can be the decision [20]
European Union	AI Act, Annex III, Articles 26 and 86	Oversight by people with "competence, training and authority"; an explanation [21]
Canada, federal bodies	Directive on Automated Decision-Making	Notice, "a meaningful explanation", "human involvement" [23]
United Kingdom, government	Algorithmic Transparency Recording Standard	Publish "how and why" a tool is used [24]
European Union, financial firms	Digital Operational Resilience Act, Article 28	A register, and "exit strategies" [25]

We did not open the Basel Committee's principles, the Australian, Singapore or Hong Kong prudential standards, the United Kingdom's prudential statements, the United States rules, any Gulf rule, or the text of any standard. An organisation should check each with its own advisers.

PART 05 · CONTINUED

What the rules have in common

01	**A decision about a person is taken by a person** who has the authority to change it.
02	**A score can be the decision** when the person who decides leans on it.
03	**The person affected can ask why,** and the organisation must be able to answer.
04	**The organisation stays accountable** for what a supplier's system does, and must be able to leave.

PART 06

A worked example: one insurer's operations department

Take a made-up mutual insurer: 2,300 staff, an operations department of 160, and about 9,000 cases a month. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

At intake, 420 new cases arrive in a day. The agent proposes the kind of case and the queue for each, and lists what is missing from 96. A case officer confirms or corrects each one.

02

Before a decision, an officer opens a disputed case with 140 pages on file. The agent's summary gives the facts with the page of each and four open points. It recommends nothing. The officer reads the two medical reports herself and decides.

03

After the decision, the officer records it with three reasons. The agent drafts the letter in the house template from those three reasons. The officer approves it and sends it.

04

In the quality sample, 60 closed cases are drawn for the month. The agent lists 23 points where a file does not show a step. The reviewer judges each one. The agent scores no case and no officer.

05

On a provider's report, the mail-handling provider's monthly report arrives. The agent lists two service levels missed and one figure that differs from the insurer's own records. The service manager raises all three with the provider.

Every one of these changes nothing by itself, sends nothing to a person outside, and leaves a record of who asked, which tool ran and what was decided.

PART 07

Seven questions before an agent touches a case

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the vendor's service? Where is the model? Several vendors can run on your side. Ask the next six of each of them.

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

03 **What can it change?**
List every tool that creates, approves, rejects, closes, files, deletes or edits a record. Is each one separate from the read tools, and does each wait for a person?

04 **What can reach a person outside?**
Show us every path by which a letter, a notice or a message leaves. Is each held for a named officer?

05 **Does it score, rank or recommend?**
Show us every score and every threshold, who set it, and whether there is a kind of case that nobody sees.

06 **What is recorded?**
Who asked, which tool ran, what was decided, for how long is that kept, and can it go into the case file?

07 **Can we stop it, and can we leave?**
One agent, one tool or one server, at once, by us. And what do we keep if the contract ends?

PART 08

Where to start

Pick work where reading is slow, the decision stays with a case officer, and nothing reaches a person outside on the agent's word.

Questions and answers on the procedures
Colleagues ask what the procedure says and get the answer with the passage cited. No case is touched.

Case file completeness check
The agent lists what a case file is missing against the checklist. A case officer asks for it.

Service provider report check
The agent lists where a provider's report departs from the contract. The service manager raises it.

One team, one kind of case, 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/OPERATIONS

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

01 Royal Commission into the Robodebt Scheme, "Report", Volume 1, 7 July 2023 (overview, and recommendation 17.1). https://robodebt.royalcommission.gov.au/system/files/2023-07/robodebt_report_volume_1.pdf

02 Federal Court of Australia, Prygodicz v Commonwealth of Australia (No 2) [2021] FCA 634, judgment of Murphy J, 11 June 2021, paragraph 4. https://www.judgments.fedcourt.gov.au/judgments/Judgments/fca/single/2021/2021fca0634

03 Post Office Horizon IT Inquiry, "Volume 1 Rapid Read" (the inquiry's own summary of Volume 1 of its final report), July 2025. https://www.postofficehorizoninquiry.org.uk/sites/default/files/2025-07/Post_Office_Horizon_Inquiry_Volume_1_Rapid_Read.pdf

04 House of Representatives of the States General (Netherlands), Childcare Allowance Parliamentary Inquiry Committee, "Unprecedented injustice" (English version of "Ongekend onrecht"), 17 December 2020, the committee's observations. https://www.houseofrepresentatives.nl/sites/default/files/atoms/files/verslag_pok_definitief-en-gb.docx.pdf

05 Autoriteit Persoonsgegevens (Dutch data protection authority), "Tax Administration fined for discriminatory and unlawful data processing", 7 December 2021. https://autoriteitpersoonsgegevens.nl/en/current/tax-administration-fined-for-discriminatory-and-unlawful-data-processing

06 UiPath, blog post "Agentic AI, now where your data lives", 20 May 2026. https://www.uipath.com/blog/product-and-updates/agentic-ai-now-where-your-data-lives

07 Automation Anywhere, blog post on a fully on-premises deployment of its platform, 16 March 2026. https://www.automationanywhere.com/company/blog/general-technology/delivering-a-fully-on-prem-enterprise-ai-automation-platform-with-nvidia-nemotron-super

08 Kore.ai, documentation page "Agent Platform overview" (deployment options), undated. https://docs.kore.ai/agent-platform ; and product page "AI for process", undated. https://www.kore.ai/ai-for-process

09 IBM, product page for its agent orchestration product, undated, read 23 November 2026. https://www.ibm.com/products/watsonx-orchestrate

10 UiPath, platform page for agentic automation, undated, read 23 November 2026. https://www.uipath.com/platform/agentic-automation

11 Camunda, "What is agentic orchestration?", last updated 11 November 2026. https://camunda.com/agentic-orchestration/

12 Automation Anywhere, "What is an Autonomous Enterprise?", undated, read 23 November 2026. https://www.automationanywhere.com/rpa/autonomous-enterprise

13 Workato, security and governance page, undated, read 23 November 2026. https://www.workato.com/platform/security

14 Microsoft Learn, "Design autonomous agent capabilities", Microsoft Copilot Studio guidance, dated 5 November 2026 in the page data. https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/autonomous-agents

15 n8n workflow templates 17806 (invoices checked against a policy and posted), 14272 (invoices from email, reviewed when a check fails), 10874 (invoices logged, then approved or rejected by a link), 19171 (IT incident investigation with an approval before the fix), 16378 (site incident reports sorted by urgency), 18513 (incident review drafts) and 7515 (chat messages sorted into tickets). https://n8n.io/workflows/17806 ; https://n8n.io/workflows/14272 ; https://n8n.io/workflows/10874 ; https://n8n.io/workflows/19171 ; https://n8n.io/workflows/16378 ; https://n8n.io/workflows/18513 ; https://n8n.io/workflows/7515

SOURCES continued

16 Eight further templates and examples: Zapier, document approval process template. https://zapier.com/templates/details/document-approval ; Zapier, template that creates a security ticket from a form. https://zapier.com/automations/it/it-service-management/it-ticketing/create-security-ticket-from-intake-form-entries-instantly ; Zapier, workplace incident reporting automations. https://zapier.com/automations/human-resources-hr/hr-compliance/workplace-incident-reporting ; Make, blog post of worked automation examples (the invoice example). https://www.make.com/blog/ai-automation-examples ; Amazon Web Services, guidance for document processing agents (a loan application example). https://docs.aws.amazon.com/solutions/intelligent-document-processing-agents-on-aws/ ; LangChain documentation, “Interrupts” (a framework pattern, not an operations template). https://docs.langchain.com/oss/python/langgraph/human-in-the-loop ; CrewAI documentation, “Human Input on Execution”. https://docs.crewai.com/learn/human-input-on-execution ; Microsoft Learn, templates overview for its agent library, dated 26 June 2026 in the page data (the one not judged). https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/kit-agent-library-templates-overview

17 150 public interface descriptions from 59 publishers, run through the platform’s importer on 23 November 2026 (Katonic research); and the search for publishers with no downloadable description (Katonic research).

18 Regulation (EU) 2016/679 (General Data Protection Regulation), Article 22(1) and (3). http://publications.europa.eu/resource/celex/32016R0679

19 Article 29 Data Protection Working Party, “Guidelines on Automated individual decision-making and Profiling for the purposes of Regulation 2016/679” (WP251 rev.01), adopted 3 October 2017, revised 6 February 2018. https://ec.europa.eu/newsroom/article29/document.cfm?action=display&doc_id=49826

20 Court of Justice of the European Union, judgment of 7 December 2023, case C-634/21, SCHUFA Holding (Scoring), paragraph 45 and the ruling. http://publications.europa.eu/resource/celex/62021CJ0634

21 Regulation (EU) 2024/1689 (AI Act), Annex III points 4(b), 5(a) and 5(d), Article 26(2) and Article 86(1). http://publications.europa.eu/resource/celex/32024R1689

22 Regulation (EU) 2026/1744 of 8 July 2026 (Official Journal 24 July 2026), the replacement of Article 113(c) of Regulation (EU) 2024/1689. http://publications.europa.eu/resource/celex/32026R1744

23 Treasury Board of Canada Secretariat, “Directive on Automated Decision-Making” (page modified 24 June 2025), requirements 6.2.1, 6.2.3 and 6.3.13. https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592

24 United Kingdom Government Digital Service, “Algorithmic Transparency Recording Standard Hub”, published 5 January 2023, last updated 8 May 2025. https://www.gov.uk/government/collections/algorithmic-transparency-recording-standard-hub

25 Regulation (EU) 2022/2554 (Digital Operational Resilience Act), Article 28(1), (3) and (8). http://publications.europa.eu/resource/celex/32022R2554

26 European Banking Authority, revised Guidelines on outsourcing arrangements. https://www.eba.europa.eu/sites/default/files/documents/10180/2551996/38c80601-f5d7-4855-8ba3-702423665479/EBA%20revised%20Guidelines%20on%20outsourcing%20arrangements.pdf

27 Financial Conduct Authority, “Operational resilience”, first published 13 May 2017, last updated 15 September 2026. https://www.fca.org.uk/firms/operational-resilience

28 International Organization for Standardization, public pages for ISO 9001:2015 (shown as withdrawn). https://www.iso.org/standard/62085.html ; ISO 9001:2026 (publication date 2026-09). https://www.iso.org/standard/88464.html ; and ISO 22301:2019. https://www.iso.org/standard/75106.html