
AI Governance for Autonomous Business Execution Human Oversight and Accountability in Practice

A framework for deciding what AI may execute, who owns the outcome and how every action stays auditable.

Written by Anne Cristea · Expert review by Samuel Bucher and Timur Göreci

WHITEPAPER 02 OCTOBER 2026 GIENI ABX


Contents

THE ABX WHITEPAPERS · 02

INTRODUCTION	The 23:40 email	03
CHAPTER 01	The AI Accountability Gap: From Principles to Practice	05
CHAPTER 02	What the EU AI Act asks of AI that acts	07
CHAPTER 03	The accountability chain	10
CHAPTER 04	Deciding what AI may do on its own	12
CHAPTER 05	When AI gets it wrong	16
CHAPTER 06	The questions your CISO, lawyer, works council and procurement team will ask	18
CHAPTER 07	Starting with governance first	24
CLOSING	The 23:40 email, answered	26
SOURCES		29

A framework for deciding what AI may execute, who owns the outcome and how every action stays auditable.

The 23:40 email



It's 23:40 on a Tuesday. An AI system connected to your ERP notices that a component is running short and emails a supplier to bring a delivery forward by a week. The email is polite, correct in its details and signed with your company's name. Nobody in your company wrote it.

At 08:15 the next morning, the supplier's account manager calls your procurement lead. The earlier delivery comes with an express surcharge, and they'd like to confirm it. Your procurement lead has never seen the email. Neither has anyone else.

The scenario is illustrative, but enterprises are starting to face it.

01

**When that call comes in,
the first question is who
approved it.**

02

**The second is where
the record is.**

In many companies running AI today, both questions have awkward answers. **This whitepaper is about those two questions.**

AI is moving from drafting text for people to carrying out work in company systems: updating records, sending messages, changing data in the ERP, placing orders. Once a system can act, the governance question changes. You need to know what it did and who allowed it, and you need to be able to prove both.

Most organisations have already written principles for responsible AI. Very few have turned them into practice. The World Economic Forum reports that fewer than 1% of organisations have operationalised comprehensive, end-to-end responsible AI practices [1][2]. The gap between the policy and the workflow is where the 23:40 email lives.

The chapters that follow give you a way to close it: a working framework for AI that acts. It covers who owns each action and how to decide what AI may do on its own, what regulators expect wherever you operate, and what happens when it gets something wrong. Read it before your next meeting with your CISO, because both questions will come up.

The gap between the policy and the workflow is where the 23:40 email lives.

The AI Accountability Gap: From Principles to Practice

Responsible AI has had a good few years on paper. Most large organisations have a policy, many have an AI committee, and some have published their principles on their website. The practice hasn't kept up.

The numbers

The World Economic Forum's AI Governance Alliance, working with Accenture, studied how organisations put responsible AI into operation. Its finding is blunt: only 10% of organisations have effective AI governance structures in place, and fewer than 1% have operationalised comprehensive, end-to-end responsible AI practices [1][2]. In the same research, 81% of companies were still at a nascent stage of implementation [3]. The researchers point to ill-defined accountability structures and protocols as one of the barriers [3].

Put simply, almost every organisation knows what responsible AI should look like and almost none has built it into the way work gets done.

Fewer than

1%

of organisations have operationalised comprehensive, end-to-end responsible AI practices [1][2]

10%

of organisations have effective AI governance structures in place [1][2]

81%

of companies were still at a nascent stage of implementation [3]

Why AI that acts raises the stakes

For the first wave of generative AI, the gap was tolerable. An assistant that drafts a summary or suggests a reply produces text, and a person decides what to do with it. If the draft is wrong, the person catches it or doesn't send it. The human is the control point by default.

That default disappears when AI starts to act. Gartner expects 15% of day-to-day work decisions to be taken autonomously by agentic AI by 2028, up from 0% in 2024 [4]. Every one of those decisions is an action somebody will have to answer for. Gartner also expects over 40% of agentic AI projects to be cancelled by the end of 2027, and names inadequate risk controls as one of the three reasons, alongside escalating costs and unclear business value [4].

So governance decides which AI projects survive. The organisations that can show who owns each action, and prove what happened, will keep their projects running. The others will pause them the first time something goes wrong and nobody can explain it.

The accountability gap

The problem has a name in this whitepaper: the accountability gap. It opens whenever AI acts inside company systems and nobody clearly owns what it does.

It rarely starts as a crisis. A wrong email to a supplier turns into a disputed surcharge. The dispute turns into an escalation. The escalation reaches an auditor, who asks for the record of who approved the action, and there isn't one. Each step is small, and each one makes the next harder to fix.



The gap keeps opening for reasons that sit outside any one team. AI tools arrive faster than governance teams can review them, and departments adopt their own. Vendors ship new capabilities every month, and each one changes what the tool can touch.

Organisations have tried to close it in familiar ways. They've published a responsible AI policy on the intranet, set up an ethics board that meets once a quarter, added monitoring after a tool went live and written separate approval rules for every workflow in every tool. Each of these has value. None of them sits where the action happens, which is why the owner and the record are missing at the moment they matter.


The rest of this whitepaper is about moving governance to that point: into the execution itself.

What the EU AI Act asks of AI that acts

The EU AI Act is one of the most detailed AI regulations in force, and it applies to many companies outside the EU that sell into the EU or use AI there. It’s worth understanding even if most of your AI use will never fall into its strictest category.

What applies, and when

FIGURE 1 · WHAT APPLIES, AND WHEN



NOT LEGAL ADVICE

DATE	WHAT APPLIES	SOURCE
2 February 2025	Prohibited AI practices and the AI literacy provisions	[5][13]
2 August 2025	Obligations for general-purpose AI models and the governance framework	[5][13]
27 July 2026	The AI Omnibus enters into force, amending the Act. The AI literacy duty is softened: organisations must now “take measures to support” sufficient AI literacy	[9][16]
2 August 2026	Most of the remaining provisions, including the Article 50 transparency obligations	[9][13]
2 December 2027	Obligations for stand-alone high-risk AI systems listed in Annex III, moved from August 2026 by the Omnibus	[9][16]
2 August 2028	Obligations for high-risk AI systems embedded in regulated products (Annex I)	[9][16]

The delay to December 2027 gives organisations time. It also means 2027 budgets and plans are being written now while the deadline is in view.

Which business uses could be high-risk

Most everyday business workflows, such as updating a CRM after a call or compiling a monthly report, sit outside the high-risk categories. Some common enterprise uses sit inside them. Annex III lists, among others, AI systems used for the recruitment or selection of people, including analysing and filtering job applications and evaluating candidates. It also covers systems used to make decisions affecting promotion or termination, to allocate tasks based on individual behaviour or personal traits, or to monitor and evaluate the performance and behaviour of workers [8].

An HR team starts using an AI tool to pre-sort applications for a new role. It saves the recruiters hours each week. Six months later, during an internal review, someone asks which category the tool falls into under the AI Act. The answer is Annex III, point 4. The team now has obligations it didn't plan for, including human oversight and log retention, and it has until December 2027 to meet them.

AN ILLUSTRATION.

The practical lesson is to classify each AI use by what it does, before it goes live. A tool's product category says little about its risk, because the same tool can do low-risk or high-risk work depending on how you use it.

What the Act asks of the organisation using the system

The Act separates providers, who build AI systems, from deployers, the organisations that use them. For high-risk systems, both carry duties. Two articles matter most for anyone running AI in their operations.

Article 14 requires high-risk systems to be designed so that people can oversee them effectively while they're in use [6]. The people overseeing the system must be able to understand what it can and can't do, stay aware of the tendency to over-trust its output, interpret what it produces, decide not to use it or override it, and stop it through a "stop" button or a similar procedure [6].

Article 26 turns that into duties for the deployer [7]. The organisation must assign human oversight to people who have “the necessary competence, training and authority, as well as the necessary support”. It must keep the logs the system generates for at least six months, unless other law says otherwise. And an employer that puts a high-risk system into use at work must inform workers’ representatives and the affected workers beforehand.

Read together, the two articles describe a governance model: a named person is responsible and has the authority to stop the system, there is a record of what it did, and the people affected are told before it goes into use. Those requirements apply by law only to high-risk systems, but they are a sound model for any AI that acts.

Those requirements apply by law only to high-risk systems, but they are a sound model for any AI that acts.

Beyond the EU

Rules differ by region and keep moving. Switzerland has no dedicated AI law in force. In February 2025 the Federal Council decided to ratify the Council of Europe’s AI Convention and to adapt Swiss law sector by sector [14][15], with a consultation draft planned by the end of 2026 [15]. Within the EU, the GDPR already gives people the right not to be subject to decisions based solely on automated processing that have legal or similarly significant effects on them [12]. Other regions follow their own timetables. As one WEF contributor put it, algorithms now travel faster than the agreements that govern them [11].

The WEF’s governance work calls for agile, anticipatory governance that can adapt as the technology changes [2][10]. For a company operating across borders, that points to one conclusion. Governance built for a single regulation has to be rebuilt every time a rule changes. Governance built into how the work runs, with an owner, an approval and a record for every action, holds up wherever the rules move.

NOT LEGAL ADVICE This chapter is a summary for orientation and does not constitute legal advice.



The accountability chain

If governance has to live inside the execution, it needs a structure every action can pass through. The accountability chain is that structure. It has six links, and an AI action is only as accountable as its weakest one.

FIGURE 2 · THE ACCOUNTABILITY CHAIN



LINK	WHAT IT MEANS	WHAT FAILURE LOOKS LIKE	TEST QUESTION
1. Owner	A named person or role is responsible for the outcome of every AI workflow, and defines what the workflow is for.	"The system did it." Nobody can say who decided the AI should be allowed to act here.	For each AI workflow you run, can you name its owner in one line?
2. Rules	The owner sets what the AI may touch, what it may do there, and what it must never do.	The AI can reach every system its credentials allow, and nobody has written down the limits.	Could you hand an auditor the list of actions this workflow is allowed to take?
3. Approval	Actions above an agreed risk level wait for a designated approver, with the authority and knowledge to accept or reject them, before they happen.	Approvals exist in some tools and not others, or they are clicked through without being read.	For your riskiest action, who approves it, and what exactly do they see before they do?

LINK	WHAT IT MEANS	WHAT FAILURE LOOKS LIKE	TEST QUESTION
4. Execution	The AI carries out the work within the rules, and only after the approvals the rules require.	The approved version and the executed version differ, and nobody can tell.	<i>Is what gets approved exactly what gets executed?</i>
5. Record	Every action is logged: what ran, on what input, with what output, who approved it and when.	Logs exist but can't be tied to a person, or they're kept for weeks when an audit needs months.	<i>Could you reconstruct any AI action from last quarter in under an hour?</i>
6. Recovery	There's a known way to stop a workflow, correct what it did and learn from it.	Stopping the AI means revoking credentials in a hurry, and nobody knows what it already changed.	<i>If this workflow went wrong tonight, who would stop it, and how?</i>

The chain maps closely to the governance model in the EU AI Act. Owner and approval correspond to the oversight duties in Articles 14 and 26, record corresponds to the log retention duty, and recovery corresponds to the ability to override and stop the system [6][7]. An organisation that builds the chain for every AI workflow is preparing for the high-risk rules even where they don't formally apply.

The chain also gives the governance conversation a shared vocabulary. When a CISO asks whether a workflow is safe, the answer stops being a matter of opinion. You can walk through six links and show where each one is covered.

✓ ABX IN PRACTICE

In Gieni ABX, the chain is built into the execution layer. Every app connected to ABX exposes the functions ABX can call, and for each one the organisation chooses whether it runs automatically, pauses for approval or is blocked entirely. Defaults are set per category and can be overridden per function. Before a change reaches a live system, ABX runs it in a sandbox and produces a reviewable difference: which fields will change, which records will be created, which messages will be sent. A person approves that exact difference, and every execution is logged with its input, output, approver and time.




Deciding what AI may do on its own

Autonomy isn't all or nothing. The same AI system can safely compile a weekly report without anyone watching, and should never change a supplier's payment terms without a person signing off. The work of governance is deciding, action by action, which level applies.

Four levels of autonomy

FIGURE 3 · FOUR LEVELS OF AUTONOMY



LEVEL	WHAT HAPPENS	TYPICAL USE
Runs on its own (human out of the loop)	The AI executes within the defined rules and records the action. No one reviews it in real time. The owner checks the record afterwards.	Internal, reversible work such as summaries and internal reports.
Human on the loop	The AI executes on its own while a designated person monitors it and can intervene or stop it at any point.	Higher-volume work where approving each action isn't practical but mistakes need catching fast.
Human in the loop	The AI prepares the action and waits. The designated approver reviews the exact action before it happens.	Anything that leaves the company or changes a system of record, including price and commission changes.
Blocked	The AI can't take this action at all.	Actions your organisation has decided should stay human, or that would bring a use into a regulated category you haven't prepared for.

A system of record is any system your organisation treats as the source of truth: the CRM, the ERP, the HR system, the finance ledger. Changes there spread to every report and decision built on them, which is why they deserve a person's approval.


Four questions that set the level

You can place any AI action by asking four questions about it.

- 1 Does the action leave the company?** An email to a customer, a message to a supplier, an order with a vendor.
- 2 Does it change a system of record?**
- 3 Can it be undone completely,** including its effects on other people?
- 4 Does it affect someone's employment or rights?** Hiring, task allocation, performance evaluation, access to services.

If the answer to the first two is no and the third is yes, the action can run on its own. If it's reversible but needs watching, a human on the loop is enough. If the action leaves the company or changes a system of record, it needs a human in the loop. If it can't be undone, or it touches someone's employment or rights, it needs a human in the loop at minimum, and may need to be blocked. And any action that answers yes to the fourth question should be checked against Annex III of the EU AI Act before it goes live [8].

FIGURE 4 · FOUR QUESTIONS THAT SET THE LEVEL



An autonomy map in practice

ACTION	LEAVES THE COMPANY?	CHANGES A SYSTEM OF RECORD?	FULLY REVERSIBLE?	AFFECTS EMPLOYMENT OR RIGHTS?	LEVEL
Compile the weekly pipeline report for the sales team	No	No	Yes	No	Runs on its own
Tag and route incoming support tickets	No	No	Yes	No	Human on the loop
Update the CRM record and deal stage after a call	No	Yes	Yes	No	Human in the loop
Send a follow-up email to a customer	Yes	No	No	No	Human in the loop
Ask a supplier to bring a delivery forward	Yes	Yes	No	No	Human in the loop
Change commission rates in the ERP	No	Yes	Partly	Yes	Human in the loop
Rank job applicants for a shortlist	No	Yes	No	Yes	Blocked until reviewed against Annex III

The map above is a starting point. Your own answers depend on your processes, your sector and the rules you operate under. The value is in having the map at all: a single document that shows, for every AI action you run, what it's allowed to do and why.

Once it exists, the conversation with your CISO changes. You arrive with every AI action already classified, the owner named and the level justified. The worksheet at the end of this chapter is designed for exactly that meeting.



ABX IN PRACTICE

In Gieni ABX, these levels are settings. Internal work that stays inside the company, such as a summary or a weekly report, can run without interruption. Anything that writes to another system or leaves the company, like sending an email, updating the CRM, changing data in the ERP or placing an order, is prepared by ABX and waits for approval. Autonomy is granted by the organisation, function by function, and never assumed.

Worksheet: your autonomy map



When AI gets it wrong

AI will get things wrong. So will people, and so will every system an enterprise has ever run. A governance model that assumes perfection breaks the first time it meets reality, and a vendor that promises perfection gives you a reason to doubt everything else it says.

The useful question is what happens next. A well-governed AI workflow handles an error in four steps.

FIGURE 5 · DETECT, CONTAIN, TRACE, CORRECT



STEP	WHAT IT TAKES
Detect	Someone or something notices. Approvals catch errors before execution. Reviews of the record catch them afterwards. Monitoring catches patterns, such as a sudden rise in actions of one type.
Contain	The workflow can be paused at once, by a named person, without revoking access across the whole system.
Trace	The record shows exactly which step went wrong: the input, the rule applied, the output, the approval. You can tell whether the error came from the data, the rules or the AI.
Correct	The effects are reversed where possible, the affected people are informed, and the rule or the autonomy level is changed so the same error is caught next time.

The last step matters most over time. Every corrected error should leave the workflow better governed than before. Sometimes that means moving the action to a stricter level. Sometimes it means a clearer rule or a better approval screen. The record of errors and corrections becomes evidence in its own right: proof that governance works when it's tested.

Being open about this builds trust with the people who have to sign off. A CISO who hears "it will never go wrong" starts looking for the hidden risk. A CISO who hears how errors are caught, contained, traced and corrected can evaluate the answer.

Every corrected error should leave the workflow better governed than before.

**ABX IN PRACTICE**

Because every Gieni ABX execution is logged step by step, an error can be traced to the exact step, input and approval involved.



The questions your CISO, lawyer, works council and procurement team will ask

Every AI initiative that acts in company systems will meet the same gatekeepers. This chapter gathers the questions they ask most often and what a good answer contains. It's written to be forwarded.

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

For the CISO

QUESTION	WHAT A GOOD ANSWER CONTAINS
What can the AI access?	Only the systems and functions explicitly connected, with the list available in writing. Access is granted per function and can be withdrawn per function.
Does our data leave the company?	The country and hosting provider where data is processed and stored. Whether your data is used to train models, and on what terms. The deployment options, including one that keeps data inside your own infrastructure. Documents your security team can check, such as a data processing agreement and an independent security certification.
Does the system remember our data?	Whether anything is kept between runs, such as past inputs, outputs or learned preferences. If so, what is kept, where and for how long. You can see it, switch it off and delete it on request. Your data stays separate from other customers' and isn't used to improve the vendor's models without your agreement.

For the CISO

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

QUESTION

WHAT A GOOD ANSWER CONTAINS

What is logged, and for how long?

Every action with its input, output, approver and time, including rejected and blocked actions. A record that can't be edited afterwards, and that you can search and export for an audit. A fixed retention period you can extend to match your audit needs. For high-risk uses, at least six months, as the EU AI Act requires for deployers [7].

How do we stop it?

A designated person with the authority to pause any single workflow immediately, without waiting for the vendor or a ticket. The other workflows keep running. Actions waiting for approval are held, not executed. The pause is recorded with who did it and when. What happens to actions in progress, and how a paused workflow restarts safely.

For Legal and the DPO

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

QUESTION

WHAT A GOOD ANSWER CONTAINS

Who is accountable for an AI action?

The organisation remains accountable, and says so internally. Each workflow has a named owner who answers for its outcome, and each action that needs approval has a designated approver, so no action sits with “the AI” or “the vendor”. The record shows both names against every action. What the contract with the vendor says about responsibility for errors, so the legal position is known before go-live.

Does any of our use fall under the EU AI Act’s high-risk rules?

Each AI use classified by what it does, not by the tool’s label, against the Annex III categories [8]. Employment uses flagged first, such as ranking applicants, allocating tasks or monitoring workers. The classification written down, dated and signed off by a named person, with Legal involved. Done before go-live and reviewed whenever the use changes. For any use that is high-risk, the date its obligations apply [9][16] and the extra duties that follow, such as oversight, logging and informing workers’ representatives.

For Legal and the DPO

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

QUESTION

WHAT A GOOD ANSWER CONTAINS

Are we making automated decisions about people?

A list of every action that affects an individual's job, pay, credit, contract or access to a service, such as rejecting an application or changing a commission rate. For each one, a human who reviews the decision and has the authority to change it, rather than approving it by default. No decision that has a legal or similarly significant effect on a person made solely by automated means, unless GDPR Article 22 allows it and safeguards are in place [12]. A way for the person affected to ask for human review and contest the outcome.

Can we show a regulator what happened?

A record for each action showing who owned the workflow, which rules applied, who approved and what executed, in one place and in the order it happened. The rules shown are the version in force at that moment, not today's. The record can be produced within days, as a file the regulator can read without the vendor's help. It holds up as evidence because it can't be edited after the event. A rehearsal: someone in the company has picked a real action and rebuilt it from the record.

For the works council

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

QUESTION	WHAT A GOOD ANSWER CONTAINS
Will this system evaluate or monitor employees?	A written yes or no, not “not by design”. If the answer is yes, or the system could be used that way, what it records about people and which uses are ruled out, such as scoring performance or tracking activity. The employee data it can reach, such as emails, calendars and activity logs, and which of it is excluded. Any use that touches worker management treated as potentially high-risk and classified before go-live [8]. Workers’ representatives and affected employees informed before the system goes into use [7], and a check whether local law requires consultation or agreement, not only information.
What happens to the people doing this work today?	The tasks that move to AI, named one by one, and the tasks that stay with people. What people do instead, such as reviewing, approving and handling exceptions, and what training or time they get to do it. Who decides how roles change, and when the people affected hear about it, before the system goes live rather than after. Whether any role, hours or pay is expected to change, and what the plan is if so. A route for staff to raise concerns, and a review after a few months to check what actually changed.
Who keeps control?	A designated approver for every action that needs approval. The autonomy map showing, action by action, which of the four levels applies: runs on its own, human on the loop, human in the loop or blocked. Control that stays with the organisation, with no action left to the AI or the vendor by default.

For Procurement

CISO

LEGAL & DPO

WORKS COUNCIL

PROCUREMENT

QUESTION

WHAT A GOOD ANSWER CONTAINS

Can the vendor show the accountability chain working?

A live run of one real task, with no slides standing in for it. The owner's rules visible before it starts. The action held for approval, with the approver seeing exactly what will change before it happens. The execution itself, then the record showing owner, rules, approver, input, output and time. A rejected action and a stopped workflow, as well as the one that succeeds. The same record produced afterwards as a file you can keep.

Does governance depend on how each workflow is built?

Governance that applies to every action by default, whoever builds the workflow. Approvals, logging and a stop control that a new workflow gets without the builder adding them. Rules set in one place and changed for all workflows at once. A test you can run: build a workflow with no governance steps and check the controls still apply.


Starting with governance first

07

Governance is what earns AI more room. The organisations that can show who owns each action and prove what happened are the ones that get permission to give AI more room.

Where to start

FIGURE 6 · STARTING WITH GOVERNANCE FIRST



- 1 Pick one workflow.** Choose one that runs often, costs your team real time and has errors you can recover from. If you're not sure which one fits, the readiness score in Whitepaper 01, "Autonomous Business Execution: Definition, Benefits and Applications", ranks your workflows so you can start with the best candidate.
- 2 Name its owner.** One person or role, written down.
- 3 Map its actions.** Run each one through the four questions in Chapter 4 and set its autonomy level.
- 4 Agree the approvers and the record.** Decide who approves what, what they see before they approve and how long the record is kept. Choose approvers for their authority and knowledge of the process, not their availability.
- 5 Run it and review the record.** After a few weeks, look at what the AI did, what was approved and what went wrong.
- 6 Widen autonomy where the record earns it.** An action that has been approved unchanged many times may be ready for a lower level of oversight. An action that keeps needing corrections stays where it is.

The trade-off

Governance has a cost, and it's better to name it upfront.

—

Actions that leave the company or change a system of record gain an approval step.

The controls apply to every workflow automatically, but each new one needs a short set-up: its owner names the actions and decides the level for each.


—

In exchange, you get a record you can show an auditor, a clear answer to the question “who approved this?” and a basis for widening autonomy with evidence.

That exchange compounds. Every workflow you govern this way becomes the template for the next one, using the same chain and the same levels. The second workflow takes less set-up than the first, and by the fifth the approach has become routine.

Every workflow you govern this way becomes the template for the next one.

The 23:40 email, answered



It's 23:40 on a Tuesday. The same AI system notices the same component running short. It prepares an email asking the supplier to bring the delivery forward by a week, and it flags that the supplier's terms include an express surcharge.

The workflow belongs to your procurement lead, who set its rules. Supplier emails sit at human in the loop, so the email waits. At 08:05 the procurement lead opens the approval, sees the exact email and the surcharge, and decides a two-day shift is enough. She edits the date and approves. At 08:06 the email goes out.

When the supplier calls at 08:15, there's nothing awkward to explain. The record shows who owned the workflow, what the AI proposed, what was changed, who approved it and when.

AN ILLUSTRATION.

Both questions from the first page now have answers. That's what AI governance looks like in practice. Every action has an owner who set its rules and an autonomy level that fits its risk, and the record shows what happened.

How Gieni ABX delivers this

Gieni ABX is an autonomous execution system for professional work. It plans, coordinates and executes work across the systems your organisation already runs, within the rules, permissions and approvals you define. Autonomy is granted function by function and never assumed. Every change runs in a sandbox first, a person approves what matters, and every action is recorded with its input, output, approver and time. Humans stay in command wherever judgement or approval is required. Governance in ABX is part of how every action executes, so it applies by default to every workflow you run, whoever builds it.

CONTRIBUTORS

Author: Anne Cristea

Expert review: Samuel Bucher, Timur Göreci

See it on a real task

Bring one workflow from your autonomy map. In a 20-minute walkthrough, Gieni ABX connects to your systems and runs it in a sandbox while you watch: the rules, the approval, the execution and the record, every step visible before anything reaches production.

[Book ABX Product Demo](#)

About Gieni ABX

Gieni ABX is an autonomous execution system for professional work. It plans, coordinates and executes work within clearly defined rules, permissions and approvals. Humans remain in control wherever judgement or approval is required.

Define the outcome. ABX executes the work. Part of Gieni AG.

THE ABX WHITEPAPERS

01 **Autonomous Business Execution: Definition, Benefits and Applications**

02 **AI Governance for Autonomous Business Execution:
Human Oversight and Accountability in Practice**

gieniabx.com

Available on Microsoft Marketplace

- [1] World Economic Forum and Accenture, **Advancing Responsible AI Innovation: A Playbook**, September 2025.
<https://www.weforum.org/publications/advancing-responsible-ai-innovation-a-playbook/>
- [2] World Economic Forum AI Governance Alliance, **Resilient AI Governance and Regulation**.
<https://initiatives.weforum.org/ai-global-alliance/resilientgovernance>
- [3] K. Yee Amezaga, R. Lazerson and M. Siddiqui, **Research finds 9 essential plays to govern AI responsibly**, World Economic Forum, 23 September 2025.
<https://www.weforum.org/stories/2025/09/responsible-ai-governance-innovations/>
- [4] Gartner, **Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027**, 25 June 2025.
<https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
- [5] **Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)**, Official Journal of the European Union.
<https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [6] EU AI Act, **Article 14: Human Oversight**.
<https://artificialintelligenceact.eu/article/14/>
- [7] EU AI Act, **Article 26: Obligations of Deployers of High-Risk AI Systems**.
<https://artificialintelligenceact.eu/article/26/>
- [8] EU AI Act, **Annex III: High-Risk AI Systems Referred to in Article 6(2)**.
<https://artificialintelligenceact.eu/annex/3/>
- [9] White & Case, **EU AI Omnibus enters into force, amending the AI Act**, 2026.
<https://www.whitecase.com/insight-alert/eu-ai-omnibus-enters-force-amending-ai-act>
- [10] World Economic Forum, **Governance in the Age of Generative AI: A 360° Approach for Resilient Policy and Regulation**, October 2024.
<https://www.weforum.org/publications/governance-in-the-age-of-generative-ai/>
- [11] J. Yu, **Building trust in AI through global governance**, World Economic Forum, November 2025.
<https://www.weforum.org/stories/emerging-technologies/trust-ai-global-governance/>
- [12] **Regulation (EU) 2016/679 (General Data Protection Regulation)**, Article 22, Official Journal of the European Union.
<https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- [13] **EU AI Act Implementation Timeline**, artificialintelligenceact.eu.
<https://artificialintelligenceact.eu/implementation-timeline/>
- [14] Swiss Federal Council, **AI regulation: Federal Council to ratify Council of Europe Convention**, 12 February 2025.
<https://www.admin.ch/gov/en/start/documentation/media-releases.msg-id-104110.html>
- [15] CMS, **AI laws and regulations in Switzerland**, updated 17 February 2026.
<https://cms.law/en/int/expert-guides/ai-regulation-scanner/switzerland>
- [16] **Regulation (EU) 2026/1744 of 8 July 2026 (Digital Omnibus on AI)**, amending Regulation (EU) 2024/1689, Official Journal of the European Union.
https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202601744

THE ABX WHITEPAPERS

- 01 Autonomous Business Execution: Definition, Benefits and Applications
- 02 AI Governance for Autonomous Business Execution:
Human Oversight and Accountability in Practice

Gieni ABX is an autonomous execution system for professional work. It plans, coordinates and executes work within clearly defined rules, permissions and approvals. Humans remain in control wherever judgement or approval is required. Define the outcome. ABX executes the work. Part of Gieni AG.

See it on a real task

[Book ABX Product Demo](#)