

ANTON by openEXPERT

The ANTON family whitepapers · Autumn 2026 · Paper 4 of 11

Trust the Process

Why confidence in AI comes from the process around it, not from chatting — and what a governed process looks like

> *Nothing in a good process asks you to trust the model. It asks you to trust the process — and it lets you check it.*

Version 1.0 · 5 November 2026

Created by Daniel Bardun · FutureChain AB

Open-source ANTON is EUPL · github.com/altspace-hub/anton_desktop

See it live ai-portal.futurechain.eu (showcases with fictional data)

Before we start

Where we left off

Last week we argued against the one big button. Real compliance work needs experts that sit close to the people who use them — with the right defaults, the right words and the right flow — and a generalist for the long tail of everything else. We ended with a thread: five experts passing one regulatory change from hand to hand, with a person approving every hand-off.

That paper was about *who* does the work. This one is about *how* an organisation comes to rely on it. In the language of the first paper, these are the trust and governance gaps.

Because here is the uncomfortable truth. You can have the best expert in the world, and if it is used in a side window — by one person, on their own, with no record — the organisation still cannot stand behind what it produces. The quality of the answer is not the issue. The issue is that nobody else can see how the answer came to be, and nobody can show that anyone checked it.

"Be careful with AI"

I have heard the same sentence from managers in many organisations over the last two years. *"Be careful with AI."* It sounds like guidance. It is not. It tells people neither what to do nor what to avoid, and it gives nobody any way to check whether they did it.

In practice it usually means something quieter: *use it, but don't tell me*. People who have seen what AI can do keep using it, on their own, outside any structure. The manager who said "be careful" has no idea how much AI is in the work that crosses their desk. And when something goes wrong, nobody can say why.

That is the opposite of trust. It is not even distrust. It is simply not knowing.

On our AI portal there is a line I wrote when we were trying to explain what the family is for, and it still says it best: *"Have trust in the process, and know that the output is going to be great. The effort goes into the set-up, not into fixing the answer afterwards."* This paper is about what that sentence means in a regulated firm — and why I believe it.

In this paper

- **Part 1, the problem:** how AI is really used today, three scenes, and why better answers do not fix it.
- **Part 2, our thesis:** confidence comes from process, the way it already does for audit — and a small test that points the same way.
- **Part 3, our take:** eight principles, what we have built, and a one-page note on AI governance.
- **Part 4, a worked example:** a client policy review, the governed way, from terms to board pack.
- **Part 5:** what this means for you on Monday, and next week.

Part 1 · The problem

How AI is really used today

Strip away the pilots and the strategy slides, and much of the AI use in compliance today looks like this.

A person opens a chat. They paste something in — a policy, a set of notes, part of a case file. They get a draft back. They edit it, sometimes heavily, sometimes hardly at all. Then they send it on, and it becomes part of the work. A paragraph in a report. A section of a policy. A summary in a weekly note to management.

Nobody else saw the prompt. Nobody knows which model it was, or which version of the policy it read. Nobody checked whether the quotes in the draft actually appear in the source. The draft became a document, and the document became the basis for a decision.

When it goes well — and it often does — nobody notices. When it goes wrong, nobody can say why, because there is nothing to look at. The work happened in a window that closed.

This is not a technology problem. It is a process problem. The model may be excellent. But the work has been lifted out of the structure that every other piece of compliance work sits inside: someone prepares, someone else reviews, the evidence is kept, and someone signs.

Three good intentions, no process

These are composites. No names, no clients. Each of them is a version of something I have seen, or heard about from people I trust.

Scene 1: the shadow draft

A team member uses a personal AI account to summarise an internal investigation for the weekly report to management. The summary is good — clear, short, well judged. It also contains customer details that should never have left the firm.

Nobody intended harm. The colleague was trying to save an hour on a busy Friday, and the tool was right there. There was simply no process, no boundary and no record. Nobody had said, in writing, which data may go to which AI service.

The lesson is not that the colleague was careless. **The lesson is that the organisation had left a careful person with no safe way to do a sensible thing.**

Scene 2: the report the client had never seen

An outsourced compliance desk sends a client its final review report. The report contains a critical finding about a control.

The client's head of finance reads it and points out, correctly, that one fact is wrong: the control was changed in June. The finding rests on the old version. But the report is already signed. It has to be withdrawn and reissued, with an awkward letter, and a little of the trust between the two firms goes with it.

Nobody did anything unusual. The analysis was careful; the reviewer was experienced. What was missing was a step: a factual-accuracy round with the client before signing — what we call in Swedish *sakgranskning* — in which the client corrects facts, not opinions, in a marked draft. One short round would have caught it.

Scene 3: the four-eyes check that was one pair of eyes

A policy is recorded in the system as "reviewed and approved". An internal auditor opens the record. Both fields show the same name: the person who drafted the policy also approved it.

Nobody meant to break the rule. The reviewer was on leave; the deadline was close; everyone assumed someone else had checked. And the tool allowed it. It may have shown a small warning. It did not refuse.

> **A four-eyes rule the system does not enforce is a wish, not a control.**

Why better answers do not fix it

It is tempting to believe that all of this will fade as models improve. It will not. Look at the three scenes again and ask what a better model would have changed.

A better model does not stop a person pasting customer data into the wrong tool. If anything, a better model makes it more tempting. The shadow draft in Scene 1 was a *good* summary. That is exactly why the colleague used it.

A better model does not make the preparer and the approver two different people. Scene 3 had nothing to do with the quality of the policy. It was about who signed it.

A better model does not give the client a chance to correct the facts before the report is signed. The fact in Scene 2 — the control changed in June — was not in any document the analysis had seen. No model, however clever, can know what nobody gave it. Only the client knew, and the process never asked.

A better model does not leave a record a supervisor can follow. A brilliant answer from a window that has closed is, from an auditor's point of view, indistinguishable from a lucky one.

These are process problems. They need process answers. And the good news — the reason I find this subject hopeful rather than gloomy — is that regulated work already knows how to build processes that earn trust. We have been doing it for decades. We just have not yet put AI inside one.

Two ways to get it wrong

An uneasy organisation tends to fall into one of two failure modes, and both are expensive.

The first is prohibition, or something close to it. On paper, the risk is managed. In practice, the people who have seen what AI can do keep using it quietly — Scene 1, multiplied — and the organisation loses the time it could have won back, just as the workload peaks with the AMLR applying from 10 July 2027.

The second is the rubber stamp. The drafts look professional, and over time people stop reading them closely. The output is trusted because it *looks* trustworthy. That is output-based trust, and it holds until the day it does not.

Between the two lies a third way, and it is not a compromise between them. **It is to treat AI-assisted work exactly as we treat any other professional work: inside a process, with roles, evidence, checks and sign-off.** Not more suspicious than human work. Not less. The same.

The draft that quietly became a decision

One more pattern sits underneath all three scenes. In a good compliance process there is a visible line between a *proposal* and a *decision*. A junior analyst proposes a rating; a senior colleague decides it. The line is visible because it is attached to named people.

In an ad hoc chat, that line disappears. The AI proposes — fluently, with confidence — and the proposal slides into the document unmarked. Nobody decided that the summary was right; it simply became the summary. Nobody decided that the finding was critical; the draft said so, and the draft was sent. **The proposal became the decision without anyone deciding.**

That, more than any error in any answer, is why managers are uneasy. And the fix is not to stop the AI from proposing. It is to make sure every proposal passes through a named person before it becomes anything more.

Part 2 · Our thesis

Trust the process, and you can trust the output.

The thesis

Confidence in AI-assisted work comes from the process around it: who does what, what evidence is kept, which checks run, who signs, and what is recorded. When the process is right, an organisation can rely on the output **without trusting the model blindly.**

That last phrase matters. We are not asking anyone to trust a model. Models change; they are updated, replaced and retired, sometimes every few months. A trust that rests on a particular model is a trust with an expiry date. A trust that rests on the process survives every change of model — because the process does not care which model proposed the finding. It cares that the finding was checked, that a person who did not prepare it decided it, and that everything is on the record.

Why we believe it

1. It is how regulated work already earns trust. Think about an audit report. Nobody trusts it because the auditor is clever, although they usually are. They trust it because of the structure around it: the auditor is independent of the work being audited; there are working papers that show the evidence; there is a review by a second, more senior person; the client is given a chance to correct the facts before the report is final; and someone signs, taking personal responsibility for the conclusion.

Every one of those elements exists because, long ago, someone learned the hard way that cleverness is not enough. **AI should sit inside that same structure, not beside it.** It is not a new kind of work that needs a new kind of trust. It is a new kind of help with an old kind of work.

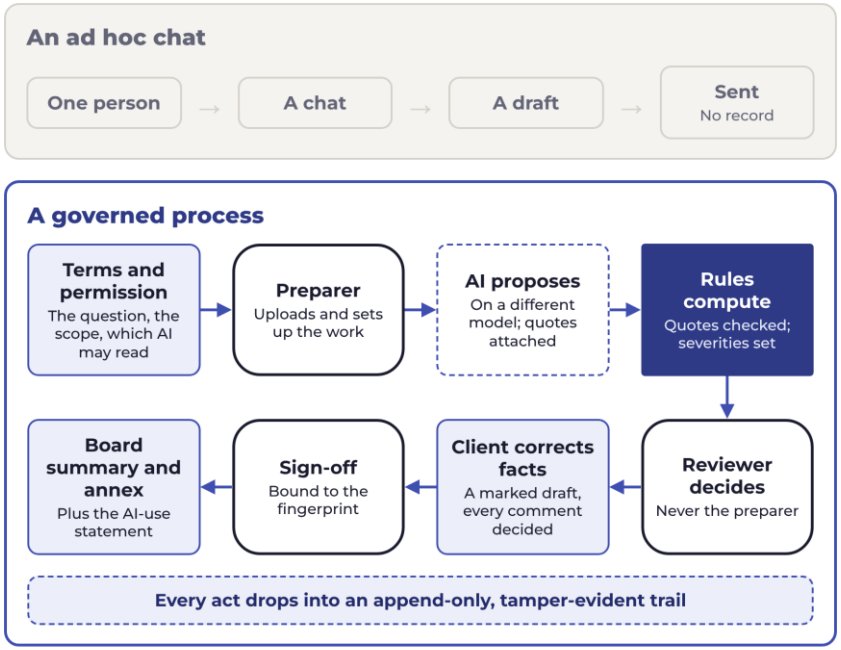
2. We said it in February, and we still believe it. Our first whitepaper, *"The Starting Point"*, put it this way: *"Process-based trust is more durable, more auditable, and more appropriate for professional and regulated contexts than output-based trust."* Output-based trust asks *"does this look right?"* Process-based trust asks *"was this done right?"* — and the second question has an answer you can check.

3. We built the process into the products, and the database enforces it. This is the reason I am most confident about. In Valentin, Olivia, Horizon and ARIS, four eyes is not a policy written in a handbook. It is a rule the database refuses to break: if the person who prepared the work tries to approve it, the system says no. In Valentin, a report *cannot* be signed while the client's factual-accuracy round is still open. In ANTON, every run is recorded and every change is audited. These are not intentions. They are properties of the software, and they hold on the busy Friday as well as on the quiet Tuesday.

4. The set-up is where the effort pays. When the method, the sources, the rules and the reader are settled before the AI writes a word, the first run is already structured and on topic. You stop fixing answers afterwards. The process is not only a safeguard at the end; it is what makes the beginning good. That is the second half of the portal line: *the effort goes into the set-up, not into fixing the answer afterwards*.

A small test that points the same way

A small internal comparison in Valentin, with models of very different sizes reviewing the same synthetic cases, pointed the same way: the set-up and the checks carried the quality, and where a model failed, the process showed the failure instead of an empty result. It is an illustration, not a benchmark, and WP6 looks at the evidence on model size carefully. But this paper's point does not rest on it. **You do not have to trust a model. You trust a process that proposes, checks, cross-checks with another model, lets a person confirm, and shows its work.**



Same AI. Different process. Different confidence.

What we are not claiming

We are **not** claiming that a process makes AI infallible. A good process makes errors visible and catchable. It does not make them impossible. That is true of every process involving people, too.

We are **not** claiming that the products make anyone "AI Act compliant", "DORA compliant" or compliant with anything else. Only the organisation can be compliant. What a governed process does is make it far easier to *show* how AI was used, by whom, under which rules — which is most of what a supervisor will ask.

We are **not** claiming that every task needs the full apparatus. A first draft of an internal memo does not need a client round and a signed opinion. The point is that the organisation should *decide* how much process each kind of work needs, and write that down — rather than leave it to whoever happens to have a chat window open.

And we are **not** using "human in the loop" as a slogan. A person in the loop means nothing unless you can say *who*, *when* and *with what power*. Throughout this paper, when we say a person decides, we mean a named person, at a defined step, with the power to accept, reject or change — and never the person who prepared the work.

Part 3 · Our take: a governed process

Eight principles

These are the principles we build to. They apply across the family, each product in its own form.

1. **AI proposes, a person decides.** Every AI output is a proposal until a named person accepts it. Nothing is recorded as done that was only proposed.
2. **Four eyes, enforced.** The person who prepared the work cannot approve it, decide its findings or sign it. That is enforced in the database, not only on the screen.
3. **Independence.** The reviewer is independent of the preparer — and the reviewing AI model is a different model from the one that wrote the document.
4. **Quality review.** The final product gets a second look before it leaves: a "second take" by another model, then a person.
5. **The client's factual-accuracy round.** The draft goes to the client to correct facts — not opinions — before signing. Their comments, and our decision on each, are recorded.
6. **Signed, frozen, fingerprinted.** The final report is frozen, given a fingerprint, and signed. Any change to the text withdraws the signature.
7. **Boundaries and caps.** Code decides what data may go to which model, with the client's signed permission. Spending is capped.
8. **A trail.** Every act is recorded in an append-only, tamper-evident trail.

Three of these deserve a word of explanation, because they are where we have made deliberate choices that cost something.

Independence of the model, not only of the person. It is now widely accepted that the person who reviews a piece of work should not be the person who wrote it. Fewer people have noticed that the same logic applies to AI. If one model drafted a document and the same model reviews it, the review shares the draft's blind spots — the same training, the same habits, the same tendency to find its own phrasing convincing. So in our review expert, Valentin, the reviewing model is never the one that wrote the document, and its verifier, by rule, is never the model that wrote the findings.

Enforced, not shown. Many tools display a warning when the preparer tries to approve their own work. A warning can be clicked away on a busy Friday — which is exactly what happened in Scene 3. We chose refusal. The database will not record the approval, whatever the screen says. **If a control can be bypassed by clicking "OK", it is advice, not a control.**

Caps as controls. Spending limits are usually thought of as a cost measure. We think of them as a control, in the ordinary compliance sense: a limit set in advance, enforced automatically, and visible when it is reached. A runaway process — a loop, a mistake, a misuse — is stopped by the cap before it becomes a story.

What we have built

Valentin: the client receives an opinion

Valentin is the family's expert for independent review, audit and model validation. Its newest release went live on the public showcase on 11 October 2026, with fictional data. It is the clearest example we have of a whole process built around AI, from the first question to the board pack.

Engagement terms come first. Before any work starts, the lead records the question being asked, who the report is for and how they may rely on it, the level of assurance (none, limited or reasonable), the scope and what is excluded, and an independence statement. Once a sign-off binds those terms, they cannot be changed. That protects everyone — the client knows what they asked for, and the reviewer cannot drift quietly into a different question.

A signed overall conclusion. For a review, the conclusion is one of four: not fit for purpose, fit with conditions, fit with comments, or fit. The database computes what the rule says the conclusion should be, given the findings that were accepted. The reviewer may depart from the rule — judgement is allowed — but only with a written reason. Whoever prepared the work cannot write the conclusion. And the text is checked, so that the conclusion cannot contradict the counts beneath it: a report cannot say "no material issues" above a list of accepted high findings.

The draft round (sakgranskning). The lead issues a draft marked *"DRAFT FOR FACTUAL-ACCURACY REVIEW"* on every page. The client's comments are recorded as received. Each comment is decided once — accepted and changed, or not accepted, with a reason. **The sign-off waits until every comment is decided.** It is not a reminder or a checklist item. The system will not let the report be signed while the round is open. Scene 2 cannot happen.

Picture: A Valentin draft export marked "DRAFT FOR FACTUAL-ACCURACY REVIEW" on every page (showcase, fictional data)

A note on language. The client does not *approve* the report: the client corrects the facts, and we sign. Practice around *sakgranskning* varies, and nothing in law requires it in this form. We built it because it is good practice — and one of the cheapest ways there is to avoid being confidently wrong.

Management's response, in two lines. For each finding, the client's management can record its response, any conditions and interim measures — briefly, on the record, beside the finding they answer.

Client packaging. With the report comes a **data-handling and AI statement**, computed from the record rather than written by hand: what was asked, where the data is kept, who is processor and controller under GDPR Article 28, how AI was used, and how outsourcing is treated under the EBA Guidelines on outsourcing arrangements. There is also a permission template for the client to sign, a data dictionary and instruction sheets.

The report split. The board gets a summary of at most two A4 pages — in Swedish unless English is asked for — with the independence statement on page one. Everything else goes into a technical annex: the fingerprints, the rules, the sign-off and the AI-use statement. Each reader gets the depth they need.

Four eyes and independence, across the family

The same idea runs through every expert, each in its own form:

- **Valentin:** the preparer cannot decide findings, confirm the scope, settle a difference between two AI opinions, clear a verifier's flag or sign. The database refuses. The reviewing model is never the author model.
- **Olivia:** reviews are done by a second person. Every invitation, password reset or access grant is confirmed by a second person. "Remediated" is kept apart from "verified". AI proposals must be confirmed by a second member before an assessment is frozen. (This last part is in Olivia's main code.)
- **ARIS:** the person who drew up a risk assessment cannot approve it, and a change cannot be approved by its owner. Approving a cycle freezes its figures and its BWRA document.
- **Horizon:** approvals need a second person with the right role, checked by the database.
- **Gabriel (the Gap Assessor):** a second opinion from a different AI model, compared exactly article by article, before a person decides.
- **ALEXANDER:** every action that would leave the system waits at a checkpoint for a person. It starts at its lowest level of independence, and only a person can raise it.

The data boundary and permission

Scene 1 was about data going where it should not. The answer is not a policy reminder. It is a boundary in code.

In Valentin, a client engagement sends **nothing** to AI until its lead has recorded the client's permission. The permission template names the AI engines the client allows, and code refuses any model the permission does not name. On the public showcase, client data has no AI route at all. Synthetic data — declared as synthetic when the engagement is created, and confirmed by a second person — may use a model pinned to one EU provider with zero data retention.

In Olivia, client data goes to AI only when both the installation and the client's lead have switched it on and recorded a legal basis and a review date. Customer-level data needs a separate legal basis, confirmed by two people. And anything that could reveal a suspicion goes to no model at all. (This is in Olivia's main code.)

Spend caps: a limit set in advance

Every product in the family caps what AI may spend. ANTON caps priced calls per day, per installation and per user — and before each call it reserves the worst-case amount, so calls already in flight count towards the cap. ALEXANDER has four caps: per run, per engagement per month, per person per day and per server per day. Valentin, Horizon, Gabriel, ARIS and the portal have daily caps. Limits set in advance, and enforced by code rather than by memory.

Trails, and ANTON's run record

In Valentin, Olivia, Horizon and ARIS, every act goes into an append-only, tamper-evident trail. In Valentin, each entry carries the fingerprint of the one before it, so a deleted or altered entry breaks the chain; in ARIS, the database itself refuses to change past entries.

ANTON, as the generalist, does the same thing in its own way. Every run writes a **run record**: the layers of the prompt with their fingerprints, the model that answered, and the tools it used. Every request that changes something writes an audit event. And seven compliance rules run after every answer. This is in ANTON's main code, and ANTON is open source, so anyone can read how it works.

> **Status at a glance**
>
> - **Live** (public showcases, fictional data): Valentin's newest release — engagement terms, the draft round, the signed conclusion, client packaging and the data boundary.
> - **In main code** (merged, not checked as deployed): Olivia's second-member confirmation and client-data rules; ANTON's run record, audit events, compliance rules and spend caps.

A short note on AI governance

This is a note, not a treatise. AI governance deserves a paper of its own. But a reader of this one will reasonably ask what the rules say, so here is the short version — what the main EU texts ask, and how a governed process helps you show it. It is not legal advice, and it is not a classification of your AI use.

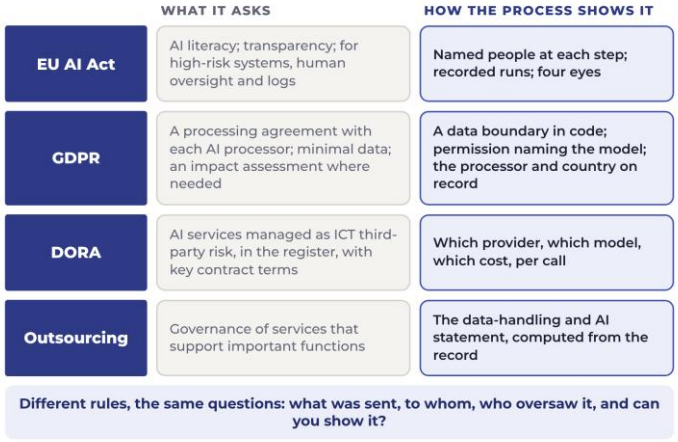
The EU AI Act (Regulation (EU) 2024/1689, amended in July 2026 by the "Digital Omnibus on AI", Regulation (EU) 2026/1744). Most compliance tooling will not be a "high-risk" AI system; the high-risk list in Annex III is specific, although it does include, for example, credit scoring of natural persons. Two duties reach almost every organisation that uses AI. Providers and deployers must take measures to support **AI literacy** among their staff (Article 4, recalibrated by the Omnibus but still a duty). And certain AI systems carry **transparency** obligations (Article 50, applying from 2 August 2026). Where a system *is* high-risk, deployers must ensure **human oversight** by competent people and keep the logs the system produces (Articles 14 and 26; logging itself is Article 12). Under the Omnibus, those high-risk obligations for Annex III systems apply from 2 December 2027. A governed process — named people, recorded runs, four eyes — is the practical way to show oversight even where it is not strictly required.

GDPR. An AI provider that processes personal data on your behalf is a **processor** under Article 28, and needs a processing agreement with the right terms. **Data minimisation** applies (Article 5(1)(c)): send only what the task needs. A **data protection impact assessment** may be needed before new AI processing that is likely to result in a high risk (Article 35).

DORA (Regulation (EU) 2022/2554, applying since 17 January 2025). For financial entities, an AI service is an ICT third-party service. It belongs in the management of ICT third-party risk and in the **register of information** (Article 28), and the contract needs the **key contractual provisions** (Article 30).

Outsourcing. Where an AI service supports an important function, the **EBA Guidelines on outsourcing arrangements** (EBA/GL/2019/02) apply to credit institutions, investment firms, payment institutions and e-money institutions.

In Sweden, the AML Act (2017:630) and Finansinspektionen's AML regulations (FFFS 2017:11) apply today. The AMLR applies from 10 July 2027.



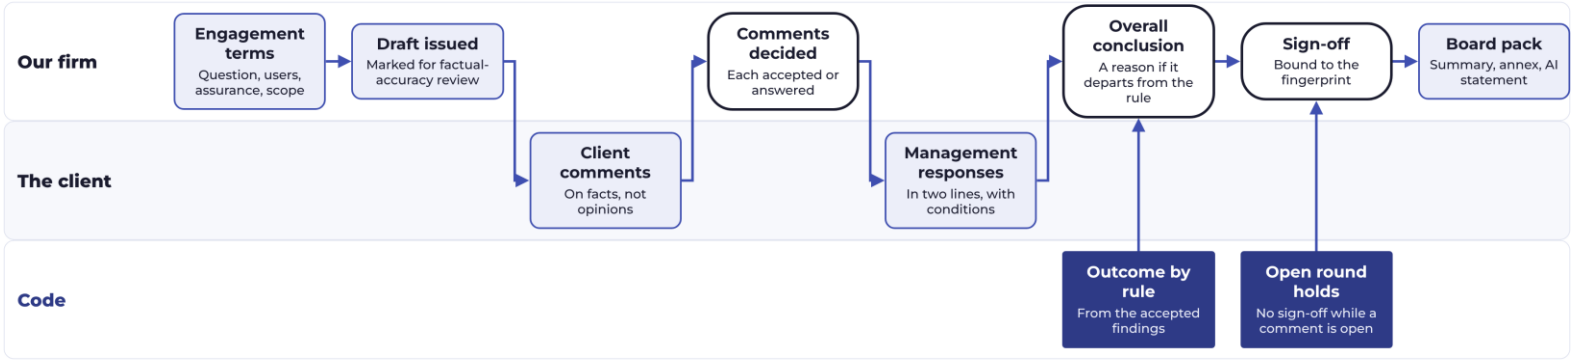
The rules ask different questions. A good record answers all of them.

What ties these together is that each of them, in its own language, asks the same few questions: *what did you send, to whom, under what agreement, who oversaw it, and can you show it?* A governed process answers each of them from the record — what went in; which model, which instructions and what it cost; who proposed, who decided and who signed; which permission named the model; which processor, in which country; and a trail that shows any change.

Part 4 · A worked example

A client policy review, the governed way

Exempelbolaget AB is a fictional Swedish payment institution. Its board has asked whether the firm's AML policy is fit to adopt. An outsourced compliance team reviews it in Valentin. Every step below runs on Valentin's public showcase today, with fictional and synthetic data standing in for a real client.



The client corrects the facts before we sign. Code makes sure.

Step	What happens	Who acts	What the record shows
1. Terms	The lead records the question (*"Is the AML policy fit for board adoption?"*), the intended users (the board), the level of assurance (limited), the scope and an independence statement.	Person (lead)	The terms, bound at sign-off
2. Permission	The client signs the permission template naming which AI may read its documents. Until then, nothing is sent.	Client, then code	The permission and the models it names
3. Preparation	A preparer uploads the policy. An upload check holds back secrets and personal identity numbers until a lead clears them.	Person and code	The upload check result
4. AI proposes	Three review lenses run on a model that did not write the policy. Each proposed finding carries a quote, and code checks every quote against the text.	AI, then code	"What went in" for each lens; the quote checks
5. People decide	A reviewer — not the preparer — accepts or rejects each finding. Rules compute the severities and the verdict.	Person, then code	Who accepted what, and the rule behind each number
6. Second take	Another model reads the draft report and flags anything that misstates the policy. A person decides each flag.	AI, then person	The flags and the decisions on them
7. Draft round	The draft goes to the client, marked on every page. The client corrects a fact: the control changed in June. The reviewer records the comment and accepts the change.	Client, then person	Each comment, and the decision on it

Step	What happens	Who acts	What the record shows
8. Conclusion and sign-off	The reviewer writes the overall conclusion, *fit with conditions*, and the rule agrees. The report is frozen and fingerprinted, and the sign-off binds to the fingerprint.	Person, then code	The conclusion, the rule, the fingerprint
9. Delivery	A two-page board summary in Swedish, a technical annex, and the data-handling and AI statement.	Code, from the record	The pack itself
10. Six months later	An internal auditor follows the whole engagement, end to end, from the trail.	Auditor	Everything above

What the auditor finds six months later

Go forward six months. An internal auditor — perhaps the one from Scene 3 — picks this engagement for a routine review. What does she find?

She finds the question the board asked, recorded before any work started, and the level of assurance the report was meant to give. She finds the client's permission, and the names of the models it allowed. She finds, for every finding, what the AI was given, what it proposed, the quote it relied on and the check that confirmed the quote. She finds the name of the reviewer who accepted each finding — and confirms it is not the preparer, because the system would not have allowed it. She finds the rule that computed each severity and the verdict. She finds the client's comment about the June change, and the decision to accept it. And she finds a signature bound to the exact text of the report, which has not changed since.

Notice what she does *not* need. She does not need to know whether the model was good. She does not need to trust the tool's vendor, or the preparer's diligence on the day. She does not need to interview anyone to reconstruct what happened. **The process is the evidence.**

That is what we mean when we say nothing in this flow asks you to trust the model. It asks you to trust the process — and it lets you check it. And Scene 2, the report the client had never seen, simply cannot occur: the system would not let the reviewer sign while the client's comment was still open.

Part 5 · From here

What this means for you on Monday

You can start building a governed process for AI with no new software at all. Five things to do this week:

1. **Move AI out of the side window.** Pick one AI-assisted deliverable and put it into your normal process: preparer, reviewer, sign-off, record. No exceptions for "it was only a draft".
2. **Make four eyes real.** Check that your tools *refuse* when the preparer approves their own work. If they only show a warning, treat the control as not enforced.
3. **Add a factual-accuracy round.** Before you sign any report about a client or a business unit, send them a marked draft to correct the facts. Record their comments and your decision on each.
4. **Write the data rule.** One page: which data may go to which AI service, under which agreement, with whose permission. Name the processor and the country.
5. **Cap the spend.** Set a daily limit per person and per team for paid AI use. A cap is a control, not a cost-cutting exercise.

Next week: the seven layers of knowledge

Next week the second act of the series begins: the method. Good AI output is mostly about context — and context is not one thing. We open the box and show seven layers of knowledge, stacked in order, that turn a general model into a domain professional: from the house rules to your own documents. And we answer one question seven times, one layer at a time, to show what each layer changes.

A closing thought

When people ask me how to make their organisation trust AI, I think they are usually asking the wrong question. Trust is not something you can install, or train into people, or win with a better demo. It is something that follows, quietly, from a process that works — one where everyone knows who does what, nothing important happens out of sight, and anyone can check.

That is how we came to trust audit reports, credit decisions and every other piece of professional work that matters. AI does not need a new kind of trust. It needs to be let into the old kind — with the same roles, the same evidence and the same signatures — and held to the same standard.

Have trust in the process, and you can trust the output. Not because the model is perfect, but because the process does not depend on it being so.

— Daniel Bardun

See it, try it, talk to us

- **See it:** the ANTON family's public showcases, with fictional data, at ai-portal.futurechain.eu.
- **Try it:** ANTON is open source — github.com/altspace-hub/anton_desktop.
- **Read the series:** all papers in the ANTON family whitepaper series, and the earlier whitepapers, are linked from the portal. [Daniel: confirm where the PDFs will live.]
- **Talk to us:** Daniel Bardun, on LinkedIn — questions, disagreements and "this is where AI stopped working for us" are all welcome.

Important notes

Product facts were checked on 11 October 2026, and the regulatory references in October 2026, including the July 2026 amendments to the AI Act; statuses are summarised in "Status at a glance". The public showcases use fictional or public data only, and no client uses them. ANTON is open source (EUPL); the other family products are not. The products support your work and help you evidence it; they do not make anyone compliant, and nothing in this paper is legal advice.