

ANTON by openEXPERT

The ANTON family whitepapers · Autumn 2026 · Paper 1 of 11

The Six Gaps

Why AI has not delivered for compliance, and what closes the gaps

> *Most compliance teams have tried AI. Most of them quietly stopped. The model was never the problem — the gap was everything around it.*

Version 1.0 · 15 October 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

A note from the author

I have spent over fourteen years in banking, financial crime prevention, regulatory work and tech consulting. In that time I have watched a long line of technologies arrive with a promise. Rules engines that would catch everything. Case management systems that would end the spreadsheet era. Screening platforms, analytics suites, dashboards for the board. Most of them worked — technically. And many of them still disappointed.

Not because the technology was weak. Because the assumptions around it were wrong. The rules engine assumed someone had written down the rules. The case management system assumed the process was clear before it was digitised. The dashboard assumed the data underneath was trusted. When those assumptions held, the technology flew. When they did not, it sat there — expensive, impressive and quietly avoided.

AI in compliance is following the same pattern. Only faster.

Over the last two years, almost every compliance team I have spoken to has tried it. The demos were genuinely impressive. The first drafts looked professional — sometimes better formatted than our own. There was a moment, for most teams, when it felt as if something fundamental had changed. And then, quietly, use dropped off. People went back to the template, the old spreadsheet, the colleague down the corridor who has done this twelve times before.

When I ask why, nobody says *"the model was stupid"*. Nobody. They say other things. *"It didn't know how we do it." "I spent longer explaining than writing." "I couldn't tell where it got that from." "Legal said no." "And who would sign it?"*

Listen to those answers carefully, because they are not complaints about intelligence. They are complaints about context, about time, about evidence, about data and about responsibility. Between them they describe six gaps that sit between a very clever model and a piece of compliance work someone can actually stand behind. None of them is closed by a bigger model. All of them can be closed — with method, with structure, and with a few firm rules about who decides what.

That is what this series is about. This first paper names the six gaps, explains why we believe they, and not the models, are the real problem, and shows what we have built to close each one. The ten papers that follow — one a week until Christmas — each go deep on one part of the answer.

One promise before we begin. We will only claim what is built and working today, and we will mark plans as plans. In a field that runs on evidence, that is the only honest way to write about AI. It is also, I think, the only way to earn the right to be listened to.

— Daniel Bardun

In this paper

- **Part 1, the problem:**
  three scenes, the six gaps, why they compound, and why this matters now.
- **Part 2, our thesis:**
  why the model is not the bottleneck, what AI is really for, and what we are *not* claiming.
- **Part 3, our take:**
  five principles and six bridges, one for each gap, and what we have built.
- **Part 4, a worked example:**
  one AMLR gap analysis, done twice.
- **Part 5, from here:**
  what this means for you on Monday, and the road to Christmas.

Who this is for

Heads of compliance, MLROs, financial crime prevention (FCP) leads, risk managers, internal auditors, and the managers they report to. No technical background is needed — the technical papers come later. If you have ever wondered why AI looked so good in the demo and so ordinary on a Tuesday afternoon, this paper is for you.

Part 1 · The problem

The quiet drop-off

Nobody announces that they have stopped using a tool. There is no meeting, no decision, no memo. It simply happens. The licence is still paid for, the icon is still on the desktop, and the work goes back to being done the way it was always done.

That quiet drop-off is the most important signal in professional AI today, and it is almost never discussed. Adoption figures measure who logged in. They do not measure who came back the following month with a real piece of work — the gap analysis due to the board, the policy that has to survive a supervisory review, the risk memo that someone will sign. That is the test that matters, and that is where AI has, so far, mostly not delivered in compliance.

To understand why, it helps to look closely at what actually happens when it fails.

Three scenes you may recognise

These are composites. No names, no clients. But if you work in compliance, you have probably lived at least one of them — or watched a colleague live it.

Scene 1: the gap analysis that was "nearly right"

A head of AML at a mid-sized payment institution asks a general AI chat to compare the firm's AML policy with the new EU AML Regulation (AMLR). Thirty seconds later she has eight tidy pages. Headings, a summary, a table. It reads well. For a moment she thinks she has just saved a week.

Then she checks it. Half the article references are vague or wrong. The output ignores the firm's own method — red, amber or green per article, with evidence, an owner and a priority — because nobody told it the method existed. And it lists rules that do not even apply to a payment institution as gaps, which means the board would have been asked to fix things the firm is not obliged to do.

She spends the afternoon fixing it, and ends up rewriting most of it. Next quarter, she writes it by hand again. Not because she is resistant to change, but because she has learned — rationally — that checking the AI's version costs more than writing her own.

Scene 2: the Monday problem

A compliance officer has become good at using AI. She keeps a document on her desktop with the paragraph she pastes in first: who we are, which regulations apply, what the board expects, how we write, what to avoid. Every session begins with twenty minutes of scene-setting before any real work happens.

Then the session ends, and it is all gone. The next week she starts from zero. Her colleague, two desks away, has a different version of the same paragraph, and gets slightly different answers. Nobody has the same setup, and nobody's setup is written down anywhere the team can share.

Setting the context takes longer than the task.

Scene 3: "who checked this?"

During a routine review, an internal auditor finds that a risk memo sent to the board was largely drafted by AI. That in itself is not a finding. But then come the questions. What was the model given? Which version of the policy did it read? Which model was it? Who reviewed the draft before it went out, and what did they change?

Nobody can answer. The memo may well be fine — it may even be excellent. Nobody can prove it.

> **Different people, different tasks, the same story. In none of the three scenes did the AI fail at writing. It failed at everything around the writing.**

The graduate who has read everything

There is a simple way to picture what is going on, and once you see it, it is hard to unsee.

Imagine you have just hired the most brilliant graduate who ever walked through your door. They have read every regulation, every guideline, every textbook and every industry report. They can analyse, summarise and write at extraordinary speed, in any language you like. On paper, they are the best hire you have ever made.

But they have never worked a day in compliance. They do not know how your firm structures a gap analysis — not the textbook version, but the version your supervisor actually expects to see. They do not know that a risk memo lands differently with the board than with the audit committee. They do not know which of your policies is the current one, or that the procedure on page twelve was replaced in March. They do not know who signs what. And every morning, they arrive having forgotten everything you told them the day before.

Would you send that graduate's first draft straight to the board? Of course not. You would train them. You would give them examples of what good looks like. You would sit next to them, check their work, and gradually trust them with more. **That is not a criticism of the graduate. It is simply how professional competence is built.**

Today's AI models are that graduate. Extraordinarily capable, genuinely useful — and almost entirely untrained in the specific, practical, experiential knowledge that separates a textbook answer from a professional deliverable. The six gaps are, in the end, a precise description of what that training, supervision and trust would have to include.

The six gaps

When we took stories like the three scenes apart — and there are many more of them — the same six gaps kept appearing. They are not about intelligence. They are about everything that sits between a model and a piece of work someone can stand behind.

1. Knowledge: most professionals are not prompt engineers

The MLRO knows exactly what a good BWRA (the business-wide risk assessment) looks like. She has written six of them. She knows which risk factors the supervisor will probe, how inherent and residual risk should connect, and where the last one was weak. What she does not know is how to tell an AI all of that, in a way it will reliably follow.

So she types one sentence, gets something generic back, and concludes that AI is not for her. The expertise was there all along. It just had no way in.

And frankly, she should not have to learn prompt engineering to get it in. We do not ask auditors to write their own spreadsheet software, and we do not ask lawyers to build their own document systems. **Expecting every professional to become an AI engineer before AI becomes useful to them is not a training problem. It is a design failure.**

2. Time: setting the context takes longer than the task

"We are a Swedish payment institution, supervised by Finansinspektionen. The AMLR applies to us from July 2027. The reader is the board. Keep it to two pages. Use our risk scale, not a generic one..." Typed again in every session, by every person, every week — and slightly differently each time.

It sounds like a small thing. It is not. A tool that saves an hour of writing but costs forty minutes of explaining is a tool people stop using. And the context that gets typed is only ever a fraction of the context that matters. The rest — last quarter's decisions, the supervisor's feedback from the last inspection, the house view on a borderline issue — never makes it into the prompt at all.

3. Training: models have read everything and done nothing

A model knows what the AMLR says. It has never sat in an assessment. It does not know how an experienced assessor weighs an article, what evidence a supervisor expects to see behind a "compliant" rating, or how a board actually reads a risk report — what they skip, what they challenge, what makes them nervous.

This is the gap I saw most clearly from inside the system. **Knowing the text is not the same as knowing the work.** The difference between the two is exactly what years of practice teach a professional, and exactly what a general model has never had.

4. Trust: "how do I know the answer is good?"

In a casual conversation, this question hardly matters. In regulated work, it is everything. A conclusion you cannot trace is a conclusion you cannot use.

A gap assessment with no visible sources, no stated method and no record of who checked it cannot go to a supervisor, however good it happens to be. And most AI tools, asked how they reached a conclusion, have only one real answer: *trust us*. In compliance, that is not an answer. **Quality you cannot show is, for practical purposes, no quality at all.**

5. Safety: "where does my data go, and who can see it?"

This is not an edge case. It is the first question any serious organisation asks, and it is a deal-breaker if the answer is unclear.

Customer files, investigation notes, transaction data and internal policies cannot be pasted into a tool nobody has assessed. When legal says no, legal is right. And until the data question has a clear, written answer, the most valuable work — the work that touches real cases and real customers — stays permanently out of reach. AI ends up confined to the safe, generic tasks where it matters least.

6. Governance: individual use is easy, organisational use needs rules

One person experimenting with AI is a pilot. Fifty people using AI with no audit trail, no spending limits, no four-eyes check and no record of which model did what is something else entirely: a finding waiting to happen.

Governance is the gap that turns personal productivity into organisational capability. Without it, every good result depends on one careful individual. With it, the organisation can show — to its board, its auditors and its supervisor — how AI is used, by whom, under which rules, and with what result.

> **Six gaps, or ten?**
>
> Readers of our first whitepaper, *"The Starting Point"* (February 2026), may remember ten gaps.
>
> Its summary named six; its chapter on the problem added four more: **repeatability**, **shareability**, **flexibility** (not being locked to one model vendor) and **accessibility**. This series uses the six, because they are the ones a compliance team feels first.
>
> The other four have not gone away. They run through the papers that follow: repeatability in WP2 and WP4, accessibility in WP3 and WP6, flexibility in WP6 and WP7, and shareability in WP10.

Why the gaps compound

It would be convenient if the six gaps were independent — six separate problems you could fix one at a time. They are not. They feed each other.

The **time** gap makes people skip the **knowledge** they should have given the model, so the output is generic. Generic output fails the **trust** test, so it has to be rewritten. Because nobody trusts it, nobody dares give it real data, so the **safety** gap never gets solved and AI stays on toy tasks. And because AI only ever touches toy tasks, nobody sees the point of **governance** — until the day someone uses it on something real, with none of it in place.

That is why so many pilots stall. Closing one gap is not enough; the others pull the result back down. **The gaps have to be closed together, by design, or they are not really closed at all.** That is the single most important idea behind everything we have built.

What the gaps cost

The obvious cost is time: the afternoon spent fixing a draft, the twenty minutes of scene-setting, the review that takes longer than the writing. But the real costs are less visible, and larger.

Teams learn the wrong lesson. After two or three disappointing attempts, a team concludes that "AI isn't ready for our work". That conclusion hardens into habit, sometimes into policy, and the organisation stops looking — just as the tools are getting good enough to matter.

AI goes underground. People who have seen what AI can do do not simply stop using it. They use it quietly, on their own accounts, outside any rules. The safety and governance gaps do not close; they just become invisible.

The backlog stays where it is. And this is the cost that matters most. Every compliance team has a list of things it knows it should do and never reaches: the better control, the remediation plan, the training for the business. When AI fails to carry the routine work, that list does not get shorter. The opportunity is not just lost — it is lost quietly, which means nobody even notices.

Why this matters now

Two things are happening at the same time, and together they make the gaps urgent rather than merely interesting.

The workload is about to peak. The AMLR applies from 10 July 2027. Every obliged entity in the EU will redo its risk assessment, its policies and its controls against one directly applicable regulation. That is a wave of gap analyses, policy rewrites and BWRA updates in a short window — landing on the same people whose calendars are already full, in a market where experienced compliance professionals are already hard to find.

AI has itself become a regulated activity. The EU AI Act asks organisations that use AI to make sure their staff have sufficient AI literacy (Article 4, applying since 2 February 2025). DORA, applying since 17 January 2025, puts ICT third-party providers — AI services included — under formal risk management. And GDPR Article 28 applies to every AI provider that processes personal data on your behalf. The governance and safety gaps are no longer just good practice. Increasingly, they are obligations.

So the question has changed. It is no longer *"should we use AI?"* It is *"how do we use AI so that we can stand behind the result?"* That is a much better question — and it has an answer.

Part 2 · Our thesis

> **The model is not the bottleneck. The gaps are.**

The thesis

A model with the right setup, working inside a process with rules and a person who signs, does useful compliance work. The same model, used ad hoc in a chat window, mostly produces confident drafts that need rewriting. So the most important question in professional AI is not *"which model?"*. It is *"what surrounds the model?"*

That sounds modest. It is not. It means the advantage in professional AI does not come from whoever has access to the largest model this quarter — those are available to everyone, and next quarter there will be a better one. It comes from whoever builds the best surroundings: the method, the knowledge, the rules and the record. **And those surroundings belong to you, not to the model vendor.** They do not reset when a new model is released. They compound.

We have argued this in depth before, in *"The Context Layer"* and in our essay *"Beyond AGI: The Case for APCI"*. Much of the industry is racing towards general intelligence. Professional work needs something else: artificial *professional, contextual* intelligence. Intelligence is becoming a utility. Context is where the value lives.

Why we believe it

We hold this view for four reasons. Each one is something we have seen, not something we hope.

1. We have seen this pattern before. Over fourteen years of implementations, the story repeats with remarkable consistency. The technology usually worked. The assumptions around it did not. The projects that went well had sorted out method, data and ownership before go-live. The ones that struggled tried to buy those things inside the tool — and discovered that no tool can supply a method an organisation has never written down. AI is no different. It is simply the most capable tool we have ever been tempted to buy them inside.

2. The gaps are not about intelligence. Make a model twice as clever and it still does not know what your organisation decided last quarter. It still forgets everything when the session ends. It still cannot tell an auditor who checked its work, or guarantee that your client's data went only where your client allowed. Each of the six gaps survives a doubling of raw intelligence intact. That is the clearest possible sign that intelligence is not what is missing.

3. We built for the gaps, one by one, and watched the difference. ANTON and its family of experts are the result of closing these gaps in working software, not on slides. Every feature described in Part 3 exists because one of the six gaps got in the way of real work. And each time a gap closed, the output changed — most visibly in how much less there was to fix — with the same kind of model underneath.

4. The setup carries further than the size. An early, small internal test pointed the same way: with ANTON's setup around it, a much smaller and cheaper open model performed on a level with a larger frontier model across a set of expert modules. It is a small sample, and we will give it the care and the caveats it deserves in WP6. But it matches what we see every day.

What AI is for: a coworker, not a headcount cut

The business case for AI is usually written as headcount reduction. I have sat in the rooms where those cases are presented, and I have never — not once — seen that promise play out the way the slide said it would.

Here is why. The assumption behind the business case is that organisations have exactly the right amount of work, and AI will let them do it with fewer people. That is not how work is. **Every professional I know has a backlog** — not of busywork, but of things they genuinely want to do and never reach. The remediation plan. The better control. The training for the business. The conversation with the product team *before* the launch, not after it.

Think about an assessment. There is value in the snapshot; you learn where you stand. But be honest: the fact that you need an assessment at all means you did not already know. **The real value was never in the assessment. It comes after it** — in the plan, the fix, the change that makes things better. That is exactly the part people rarely have time for, because they are already busy with the next assessment.

AI as a coworker changes that equation. It carries the necessary, recurring work faster, so people get to the work that moves things forward. And something else happens too. When a gap analysis takes days instead of weeks, it no longer has to happen once a year. It can happen every quarter, or whenever the regulation, the business or the risk picture changes. You stop running an annual check-up and start running continuous diagnostics. **That is a change of rhythm, not of headcount** — and it is the subject of WP9.

> **AI is not here to replace your colleagues.**
> **It is here to become one.**

What we are not claiming

A thesis is only useful if you are clear about its limits, so here are ours.

We are **not** saying models do not matter. Better models reason better, and for some hard, open-ended problems a larger model still makes a real difference. WP6 is honest about where.

We are **not** saying AI makes anyone compliant. Nothing does that except the organisation itself. Good tools help you do the work and *show* the work. That is a lot — but it is not the same thing.

And we are **not** saying people step back. The opposite: in every design choice we have made, the person who knows the work stays in charge of the conclusion. AI proposes. A person decides.

Part 3 · Our take: six bridges

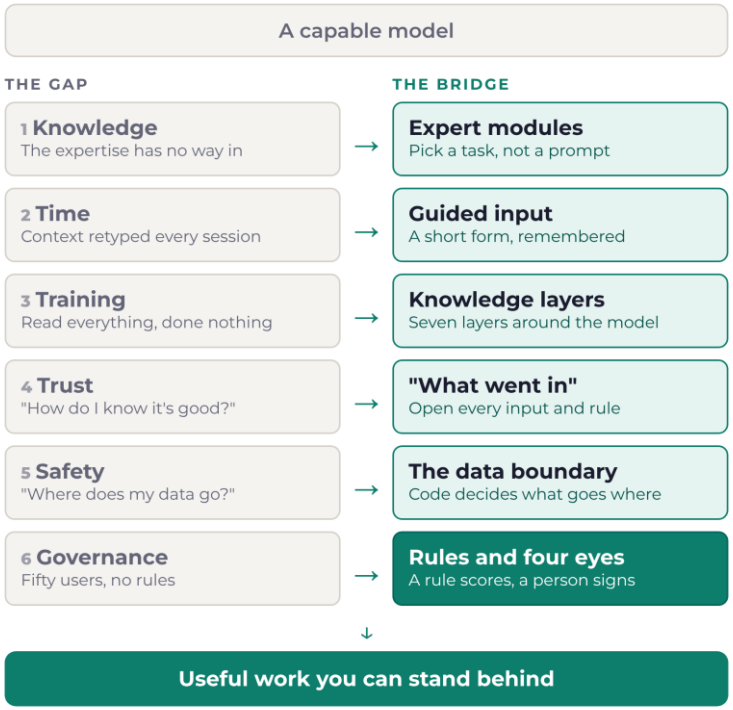
Five principles first

Before any feature, five principles. They are the reason the features look the way they do — and they are the lens for the rest of this series. If you remember nothing else from this paper, remember these.

1. **Start with the problem, not the tool.** You pick a task — "AMLR gap analysis", "review this policy before the board" — not a model and not a prompt. Professionals think in tasks and deliverables. The tool should too.
2. **Put the expertise in the setup, once.** The method, the sources, the format and the reader are settled before the AI writes a word. That is what saves the rewrite, and it is what makes the result repeatable from one person to the next.
3. **Show what went in.** Every answer can be traced to the instructions, the documents and the model that produced it. Not as a promise — as something you can open and read.
4. **Keep the data on your side.** Your documents stay with you. Only what a task needs goes to a model, under rules you set, and the record says where it went.
5. **AI proposes, code computes, a person decides.** Where a number has to stand up — a gap score, a risk rating, a review verdict — a fixed rule computes it, not the AI. A person signs, and never their own work. That is "four eyes": the person who prepared the work cannot approve it.

The fifth principle deserves a moment, because it runs against the grain of most AI products. It would be easy to let the model produce the score; it is fluent, fast and confident. But a score that comes from a model can change when the model changes, cannot be fully explained, and cannot be reproduced on demand. A score that comes from a written rule can be read, challenged, tested and reproduced by anyone, next week or in three years. **In regulated work, the second kind is the only kind worth having.**

Six gaps, six bridges



The model was never the problem. Six gaps were — and each one can be bridged.

Knowledge → expert modules. You choose a task from a catalogue, and the expert instructions are already written — the method, the structure, the questions an experienced professional would ask, the pitfalls they would avoid. ANTON has more than 550 ready-made modules in 60 areas, from AML gap analysis and the BWRA to policy drafting, legal research, risk and audit. The MLRO from the knowledge gap no longer has to translate her expertise into a prompt. Someone who knows the work has already done the translation — once, for everyone.

Time → guided input. Instead of a blank box, a short form asks for what actually matters: jurisdiction, sector, audience, documents, the output you need. Settings are remembered, so Monday does not start from zero, and two colleagues running the same task start from the same place. Twenty minutes of scene-setting becomes a minute of choosing.

Training → knowledge layers. Seven layers stack the professional training around the model, in a fixed order: house rules, the area's context, the module's method, a persona, skills, your own sources, and the depth of reasoning. More than 40 regulatory knowledge packs ship with ANTON, from the AMLR and the EBA guidelines to DORA, MiCA, GDPR and the AI Act. This is how the graduate gets their training — not by becoming a bigger graduate, but by being given what a professional would know. WP5 opens all seven layers.

Trust → "what went in". For any answer, you can open the exact instructions, sources and model behind it, and read the rule that computed every number. Our review expert, Valentin, goes a step further and publishes its whole method: every prompt as sent, and every rule as a "rule card" with its formula, its inputs and a worked example. When the auditor from Scene 3 asks *"what was the model given?"*, the answer is one click away.

Safety → the data boundary. Code, not good intentions, decides what data may go to which model. In Valentin, client data goes only to a model the client has permitted, and the record names it. ANTON is local-first: your documents stay on your own machine, and only the call to the model leaves it. On the public showcases, AI calls are pinned to an EU provider with zero data retention. The legal question gets a written answer before the work starts, not after.

Governance → rules and four eyes. Fixed rules compute the scores. The preparer cannot approve their own work — and that is enforced in the database, not in a policy document nobody reads. Every act lands in a tamper-evident trail, and spending is capped. Governance stops being a set of good intentions and becomes a property of the system.

> **Where a number has to stand up, a rule computes it. A person signs. Never their own work.**

Notice that none of the six bridges is a better model. Every one of them sits *around* the model. That is the thesis, made concrete.

What we have built

ANTON: the all-round expert

ANTON is one workspace with more than 550 ready-made expert modules in 60 areas. You pick the problem and the output you need — an executive summary, a gap matrix, an action plan — and ANTON assembles the expert instructions in seven layers. Every run is recorded, so it can be replayed and audited later. It runs on your own machine; only the call to the model leaves it — and with a local model, not even that.

ANTON is open source under the European Union Public Licence (EUPL), because we believe this kind of capability should not be reserved for the organisations that can afford the largest consulting budgets. A showcase with fictional data is live at anton.futurechain.eu.

The family: experts close to the work

ANTON can do almost anything. But where a job needs more than a good draft — where it needs its own rules, its own records and its own sign-off — a focused expert does it better:

- **Gabriel (the Gap Assessor):** article-by-article gap assessments against 62 frameworks. The AI extracts the facts; a fixed rubric gives the score and the red, amber or green rating.
- **Valentin:** independent review and a second opinion before you sign; model validation and audit; signed reports bound to the exact version that was reviewed.
- **ARIS:** the bank's business-wide risk assessment under the AMLR, with every figure traceable to a person, a rationale and a method.
- **Horizon:** regulatory horizon scanning for financial crime, in the law's own words, checked against the official text.
- **Olivia:** the workspace for outsourced compliance teams running many small clients' *årshjul* (the annual compliance cycle), each client walled off from the others.
- **ALEXANDER:** the continuous expert for a standing mandate, re-checking on a schedule or whenever something changes.

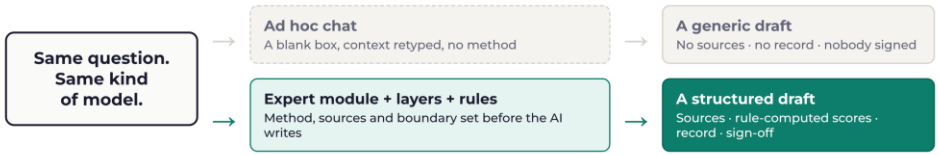
All of them are live on public showcases with fictional data, gathered in one place at ai-portal.futurechain.eu. Unlike ANTON, they are not open source. More experts are on the way, and WP3 introduces the whole family — and explains why a family of focused experts beats one big

Part 4 · A worked example

One AMLR gap analysis, done twice

Exempelbolaget AB is a fictional Swedish payment institution. Its head of compliance needs an AMLR gap analysis for the board. On the left, she does it the way most teams try first: in an ad hoc chat. On the right, she does it with the gaps closed. **The kind of model underneath is the same in both columns.** Everything else is different.

Step	In an ad hoc chat (the gaps open)	With the gaps closed (ANTON, Gabriel and Valentin)
1. Start	A blank box. *"Compare our AML policy with the AMLR."*	She picks the task: *AMLR gap analysis*. The module's method is already there. **Knowledge gap closed.**
2. Context	She types who they are, which rules apply, what format, for whom. And again next time.	A short form: entity type (payment institution), country (Sweden), audience (the board), output (gap matrix and action plan). Remembered next time. **Time gap closed.**
3. Knowledge	Whatever the model happens to remember about the AMLR.	The AMLR knowledge pack and the article text, the firm's own policy, and the house method as a skill. Articles that do not bind a payment institution are marked out of scope. **Training gap closed.**
4. Scoring	The model writes "partially compliant" wherever it feels like it.	The AI extracts facts per article: documented? implemented? tested? owner? A fixed rubric turns the facts into the score and the rating. Same facts, same score, every time. **Governance gap closed.**
5. Data	The policy is pasted into a tool nobody has assessed.	The policy stays on the firm's own machine or server. The call goes to a model the firm has approved, and the record says which. **Safety gap closed.**
6. Check	Someone skims it.	A colleague reviews it in Valentin: findings with quotes checked against the text, a second take from a different model, and a sign-off bound to the exact version. The preparer cannot sign their own work. **Trust gap closed.**
7. Afterwards	Nobody can say what went in.	"What went in" shows every instruction, source and model. The report carries its own method annex.



Same model. Different setup. Different result.

What changed, and what did not

Look at the right-hand column again and notice what is *not* there. There is no magic. No step depends on the model being cleverer than it was on the left. Every improvement comes from something that was decided before the AI wrote a word, or checked after it finished: a method, a form, a knowledge pack, a rule, a boundary, a second pair of eyes.

Notice, too, where the person is. On the left, the head of compliance is a proofreader, fixing a draft she did not shape. On the right, she is doing the work only she can do: choosing the scope, judging the evidence, deciding what to tell the board. The routine work has moved to the coworker. The judgement has stayed with the professional.

And notice the last row. On the left, the work is finished when the document is sent. On the right, the work leaves a trail — which means it can be shown to an auditor, repeated next quarter, and improved the time after that. **That is the difference between using AI and being able to rely on it.**

Part 5 · From here

What this means for you on Monday

You do not need any of our software to start closing the gaps. Most of the work is thinking, deciding and writing things down. Five things to do this week:

1. **Name your gaps.** Take one AI attempt that disappointed your team and ask which of the six gaps it fell into. Most fall into more than one — which is exactly why closing just one did not help.
2. **Write down your method, once.** Pick one recurring task, a gap analysis or a policy review, and write the half page that says how your team does it and what good looks like. That half page is the beginning of your setup — and it is valuable even if you never use AI.
3. **Ask the data question first.** Before the next pilot, write down which data may go to which AI service, under which agreement (GDPR Article 28), and who decided. If nobody can answer, that is the first gap to close.
4. **Decide who signs.** For any AI-assisted deliverable, name the person who checks it, and make sure it is not the person who prepared it.
5. **Measure the rhythm, not the headcount.** Pick one check you do once a year and ask what it would take to do it every quarter. That is where AI as a coworker pays.

The road to Christmas

This paper is the map. The next ten papers are the territory. Each one goes deep on one gap, or on what becomes possible once the gaps are closed.

The series moves in three acts.

Act 1 is the problem and how we think: why an answer you cannot explain is worth nothing, why experts close to the work beat one big button, and why trust comes from process rather than from chatting.

Act 2 is the method: the seven layers of knowledge, why setup beats size, how to keep your knowledge on your side while models come and go, and atoms — the small pieces of learning that let each task start smarter than the last.

Act 3 is the new way of working: checks that run whenever something changes, one openly documented format for sharing expertise, and ANTONs that connect safely across people, teams and organisations.

WP1 · The six gaps — 15 Oct · you are here

Act 1 — The problem and how we think

WP2	No mystery box — Trust	22 Oct
WP3	Not one big button — Knowledge	29 Oct
WP4	Trust the process — Trust · governance	5 Nov

Act 2 — The method

WP5	The seven layers of knowledge — Training	12 Nov
WP6	Setup beats size — Time · cost	19 Nov
WP7	Your knowledge, your side — Safety	26 Nov
WP8	Atoms: a learning loop — Training over time	3 Dec

Act 3 — The new way of working

WP9	From once a year to all the time — Time	10 Dec
WP10	The .anton format — Sharing	17 Dec
WP11	Connected, and 2027 — Connection	23 Dec

Eleven weeks, one story: from why AI has not delivered to how we will work in 2027.

Go deeper: the papers behind this series

- **ANTON whitepaper part 1, "The Starting Point"** (February 2026). The original gaps, and the foreword on AI as a coworker.
- **Part 2, "The Context Layer: How APCI Changes Professional AI"** (March 2026). Why intelligence is not the bottleneck.
- **"Beyond AGI: The Case for APCI".** Our essay on why the value in AI compounds in the context layer, not the model layer.
- **Part 3, "The Network: When ANTONs Connect"** (April 2026). The subject WP11 returns to.

These are longer and more technical, and some of their figures have since moved on. The numbers in this series are current as of October 2026.

Next week: no mystery box

An answer you cannot explain is worth nothing in regulated work — however good it is. Next week we look at why every AI-assisted conclusion should show what went in, how each number was computed and who confirmed it, and what a "glass box" looks like in practice. Our motto fits in three words: **explainable, reusable, understandable.**

A closing thought

The model was never the problem. That is, when you think about it, very good news. It means we do not have to wait for the next model, or the one after that, for AI to become useful in compliance. The gaps are things we can close ourselves — with method, with knowledge, with rules and with people who sign. None of it depends on a breakthrough. All of it compounds.

The organisations that do well with AI will not be the ones with the biggest model. Everyone will have access to that. They will be the ones that built the best surroundings for it — and then gave their people the time to do the work they never had time for.

And those surroundings point somewhere. Within a few years, we believe, good compliance work will run continuously rather than by the calendar; every conclusion will carry its own explanation; organisations will own their professional context while models are bought and swapped like utilities; expertise will travel, packaged and checked, between colleagues, teams and firms; and people will spend their time on judgement, improvement and the conversations that matter. That is what we mean by artificial professional contextual intelligence, and it is what we are building towards — one honest step at a time.

That is what this series is about. I hope you will follow it, argue with it, and tell us which gap stopped AI in your team.

— Daniel Bardun

About

The ANTON family whitepapers are written by Daniel Bardun, who has spent over fourteen years in banking, financial crime prevention, regulatory work and tech consulting, and published by FutureChain AB. ANTON by openEXPERT is an open-source AI workspace for professional work; the ANTON family adds focused experts for compliance tasks that must stand up in front of a supervisor.

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; 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.