

ANTON by openEXPERT

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

Connected

Why expertise should travel safely between people, teams, agents and organisations — what we have built, and what we have not yet proven

> *Signed out, signed back, read by nobody in between. And a person approving every commitment.*

Version 1.0 · 23 December 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 opened up the .anton file: a package that says what it is, proves it has not changed, and cannot run anything. We argued that expertise should travel as a package rather than as a paragraph, so that good work gets reused instead of rebuilt, and trust can be judged instead of assumed.

A package answers *what* travels. This last paper of the year is about *how* it travels — and between whom. Not only as a file attached to an email, but as a question asked of a colleague's AI, a task handed to another team's installation, or a request from one AI agent to another. In the language of the first paper, this is not one of the six gaps. It is what closing the six makes possible: the connection between people, teams and organisations.

Connecting the islands

In our April whitepaper, *"The Network"*, we described the problem with a picture I still find useful: **islands**. Every person's or team's AI set-up gets better on its own — its modules, its knowledge packs, its record of what worked. But the islands do not talk to each other.

The office that worked out a tricky beneficial-ownership question last month cannot share the answer with the office that needs it today. The same gap is found independently in three places, and nobody sees the pattern. And when the islands *do* get connected, it is usually by moving everything into one provider's cloud — which regulated firms rightly hesitate to do.

At the same time, something new is happening. AI agents — software that can take steps for you, within limits you set — are starting to act for people: asking for documents, negotiating, even paying. An agent that can commit on your behalf, without the right guardrails, is a new kind of risk.

This paper is about connecting the islands without building a centre that reads everything, and about letting agents work without letting them decide. It is also, because it is the last paper in the series, a look back at the six gaps and a look forward to 2027.

Two facts this paper keeps visible

- **Not released.** Most of what this paper describes is built and tested on real phones and computers. It is not publicly released and not in any app store.
- **No independent external security review yet.** The relay, the mesh and the agent gateways have had internal reviews only. Until an independent review is done, we recommend keeping client work on your firm's approved channels — and that includes ours.

In this paper

- **Part 1, the problem:** islands of expertise, three scenes, and why "one big cloud" is the wrong answer.
- **Part 2, our thesis:** peer to peer, selective, signed and approved by a person — and what we are *not* claiming.
- **Part 3, our take:** seven principles, what we have built, and the four uses of a connected network, each with an honest status.
- **Part 4, a worked example:** asking another ANTON's expert, from question to signed answer.
- **Part 5:** what this means for you on Monday, a look back over eleven papers, and 2027.

Part 1 · The problem

Expertise on islands

For most of this autumn we have been building up the island. The seven layers of knowledge, the expert modules, the house rules, the atoms of learning, the package that lets a team's method travel. Each of those makes one person's or one team's AI set-up better.

But compliance work has never been done on one island. An MLRO depends on the business. A second-line reviewer depends on the first line. A group function depends on its subsidiaries in other countries. An investigator depends on what a correspondent bank's compliance team is willing to tell her. **The knowledge that matters is almost always partly somewhere else.**

Today, that knowledge moves the way it always has: by email, by phone, by a meeting where someone happens to mention it. Slow, unrecorded and dependent on who knows whom. AI has made each island faster. It has done almost nothing, so far, for the water in between.

Three islands

These are composites. No names, no clients.

Scene 1: the colleague with the answer

An analyst in Stockholm needs to know how a Finnish rule on beneficial ownership differs from the Swedish one. A colleague in the Helsinki office answered exactly this question last month, carefully, in her own AI workspace. The answer exists.

But there is no safe way to ask her AI. The analyst emails the colleague, who is travelling, and waits three days for a reply that summarises, from memory, what her workspace already holds in full.

Scene 2: the counterparty request

During a customer investigation, a bank's investigator needs a document from a correspondent's compliance team. Today that means emails, attachments and phone calls. Afterwards, there is no single record of what was asked, by whom, what was given, and on what basis.

Everyone involved is acting in good faith. But if a supervisor later asks how that document was obtained, the honest answer is a forwarded email chain.

Scene 3: the agent that paid

A team trials an AI agent that can book services on its behalf. It is impressive — it finds options, compares prices and gets things done. One morning it pays for something nobody approved.

The investigation is short. "Approval" was a setting, and the agent could change its own settings. **The guardrail was inside the reach of the thing it was meant to guard.**

In each scene, the expertise or the capability existed. What was missing was a safe way for it to travel — and a person in the right place when it did.

Why "one big cloud" is the wrong answer

The obvious way to connect islands is to put them on one continent: move everyone's work into one shared platform in one provider's cloud, and let it all talk. Many tools work this way, and for many purposes it is perfectly sensible.

For regulated work it raises three problems at once.

It concentrates what should be separate. A compliance function holds information under duties of confidentiality, data protection and, in banks, secrecy. One central place that can read everything is one place whose failure exposes everything — and one more party in the chain that has to be assessed, contracted and supervised.

It takes the decision away from the owner. When everything sits in one shared space, sharing becomes the default and keeping things private becomes the effort. In regulated work it must be the other way round. **Sharing should be a choice someone makes, not a state someone forgets to change.**

It makes the provider the point of trust. In a central model you trust the platform to keep each party's data apart. That may be well founded. But it is trust in a promise, and this whole series has argued that trust in regulated work should come from what can be shown, not from what is promised.

Our April whitepaper put the principle in one line that I still hold to: the value of a network grows with its connections, **but each node must keep its sovereignty.** The answer to isolation cannot be centralisation.

What "encrypted" does and does not tell you

One more thing before the thesis, because it is where a great deal of marketing is quietly misleading.

"End-to-end encrypted" tells you something important about *content*: that the words in a message can only be read by the sender and the receiver, not by whatever carries it in between. It tells you nothing about *metadata* — who connected, to whom, when, and how much. Metadata can be revealing on its own. That a compliance team contacted a particular correspondent on the day before a report was filed may say a good deal, even if nobody can read what was said.

So the honest description of any connected system has two halves: what the carrier cannot see, and what it can. We will give both.

Part 2 · Our thesis

> **Expertise should travel safely: peer to peer, by choice, signed, and with a person approving every commitment.**

The thesis

Connection between AI set-ups should have four properties.

1. **Peer to peer and encrypted.** No central place that reads everything. A relay may carry messages, but it cannot read them.
2. **Selective and consented.** You decide exactly what is shared, with whom. Nothing is shared by default.
3. **Signed.** You know who did the work, and you can tell if anything changed on the way.
4. **Approved by a person.** No agreement is signed, no money moves and no commitment is made without a person — and the approval sits outside the AI's reach.

If a connection has all four, colleagues, teams, agents and organisations can share knowledge and work without giving up control. If it is missing any one of them, it is either a silo or a risk.

Why we believe it

1. Sovereignty and connection are not opposites. We argued this in April and nothing since has changed my view. Every node — a person, a team, a firm — keeps its own data, its own keys and its own decisions. What it shares, it chooses to share. That is not a compromise between privacy and collaboration. It is the only form of collaboration a regulated firm can actually sign up to.

2. We built the pieces and tested them on real phones and computers. Encrypted messaging between people. Contact between ANTON installations, with signed task hand-offs. A relay that forwards sealed envelopes. Local gateways where an agent can negotiate but a person must approve. None of this is a slide; all of it has run, end to end, on real devices.

3. We are honest about the limits — and the limits are part of the argument. Built and tested is not the same as reviewed. Nothing we describe here has had an independent external security review yet. I would rather say that on page one than have a reader discover it on page ten. In a field that runs on evidence, **a stated limit earns more trust than a confident claim** — and it is the same standard we have applied to every paper in this series.

An analogy: the sealed letter

Think about how post works. You write a letter, seal it, and put an address on the envelope. The postal service carries it across the country. It sees the address, the postmark and the weight. It does not open the letter. And the person who receives it can see your signature at the bottom and judge whether it is really from you.

That is the shape of what we have built. **The relay is the postal service:** it carries sealed envelopes between installations and phones, and it sees who is writing to whom and when — the address and the postmark — but not what is inside. **The signature is your signature**, checked by the receiver. And the most important part has no postal equivalent: **when the letter asks for a commitment — an agreement, a payment, a piece of work — a person opens it and decides.** Not the postman, and not the letter.

Like every analogy it has edges. The postal service is a public institution with a long record; our relay is new, and it has to earn that record.

What we are not claiming

We are **not** claiming the system is ready for enterprises or banks, or that it cannot be broken. It has had internal reviews only. The current version of the mesh is designed for households, NGOs and small firms — not for firms that must inspect their own network traffic.

We are **not** claiming that the relay "knows nothing". It sees metadata: who connects, and when. It also keeps undelivered messages for up to seven days, sealed, so they can be delivered when the receiver comes back online.

We are **not** claiming that agents exchange investigation files or transaction data between financial institutions. They do not, and that use would need a legal basis, secrecy rules and four eyes before it could be built.

We are **not** claiming that any of this is released. It is built and tested, on real devices, and not yet available to the public.

And we are **not** claiming that connection replaces judgement. An answer from a colleague's AI is a second opinion, not a fact. A person still decides what it means.

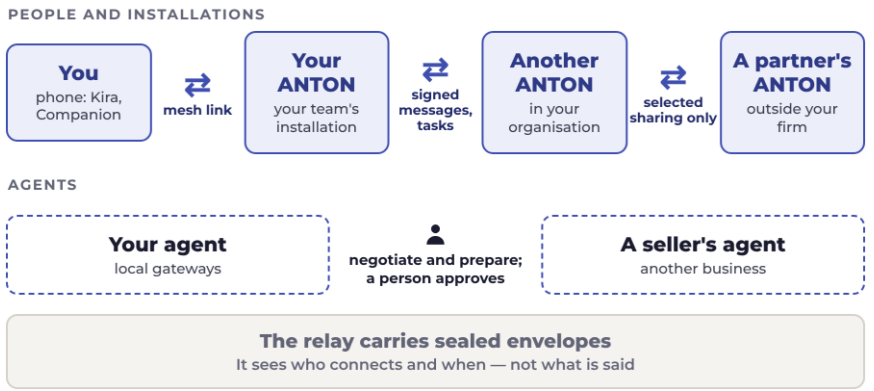
Part 3 · Our take: connected, without a centre that reads everything

Seven principles

1. **The relay passes messages; it cannot read them.** It sees who connects and when, not what is said.
2. **Keys stay on your devices.** Phone keys in the phone's secure store; installation keys on the installation.
3. **Pair in person.** People and agents are added with a pairing code or a QR code, so you know who is on the other end.
4. **Share selectively.** Only what you choose to share leaves. For knowledge, only what is marked as shared — never private project knowledge.
5. **Sign the work.** A task handed to another ANTON goes with a signed brief and comes back with a signed result. A failed signature makes the result unusable.
6. **AI proposes; a person approves.** No agreement is signed and no money moves until a person approves, outside the AI's reach.
7. **Say what has been checked.** No security claim beyond what has been reviewed.

Principle six is the one Scene 3 was about, and it is the one I care most about getting right. It is not enough for an agent to *ask* for approval. The approval has to happen somewhere the agent cannot reach. In what we built, the agent prepares the agreement or the payment, and then it stops. A person approves in a separate window — in a terminal or a browser — and **the link to that approval is never shown to the AI**. The agent cannot click it, cannot change it and cannot pretend it was clicked.

Principle three sounds old-fashioned, and it is meant to. The weakest point of most secure systems is not the encryption; it is knowing who is on the other end. Pairing with a code exchanged in person, or a QR code scanned on the spot, is a small ceremony. It is also the step that turns "encrypted" into "encrypted, with the person I think it is".



Connected, without a centre that reads everything.

What we have built

Kira: one communication layer for people, agents and ANTONs

Kira is the name of the layer that ties the pieces together. Its parts share one relay. All of them are **built and tested, and not released**.

Person to person. The Kira phone app, on Android, gives encrypted chats, group chats and calls. Messages are encrypted on the phone before they leave it, and calls go directly between phones wherever the network allows.

You to your own agent. The ANTON Agent phone app lets you give tasks to your own AI agent and read its replies. One limit belongs here: forward secrecy — which keeps old messages safe even if a key is stolen later — is built into the layer, but we do not yet claim it, and this channel between the phone app and your own agent does not have it.

Your agent to other businesses' agents. On your own computer, two small gateways give an AI agent tools. **Collaboration** is a task inbox: finding sellers, negotiating, preparing signed agreements, and keeping records of delivery and escrow. **Agent Pay** is a wallet with spending limits; the agent may propose a payment, never send one. Each gateway waits for a person to approve, in the way I described under principle six. They work with any agent that speaks MCP, an open protocol for giving AI agents tools. Together they make up the **ANTON Agent Standalone**, which needs no other part of ANTON. Today they are a local install for technically confident users.

ANTON to ANTON. Two ANTON installations connect by exchanging contact codes. They can then send each other signed messages, encrypted end to end once both have exchanged keys. They can ask each other's specialist agents a question, and hand over a piece of work — a "mission task" — with a signed brief and a signed result, and a notice when the other side accepts or declines. Today these messages go through a signed queue that delivers roughly every 30 seconds; the real-time protocol is not finished.

Your phone to your own ANTON. The **mesh** link lets your phone reach your own ANTON from any network, through the relay, without opening your computer to the internet. Both sides dial *out* to the relay, so nothing has to listen for incoming connections. It uses the same well-studied handshake design as WireGuard, a widely used encrypted-tunnel technology, and it comes with a written threat model. This first version is designed for households, NGOs and small firms that host their own ANTON — not for enterprises.

The Companion app

The Companion app pairs a phone with your desktop ANTON. Pairing uses a short-lived enrolment code and a fresh key pair created on the phone, so the phone's key never leaves it. Once paired, you can ask your ANTON a question, capture a document with the camera, and — most usefully — **approve pending actions** from wherever you are, with a biometric re-confirmation for the high-severity ones. Push notifications carry no confidential content; they tell you there is something to look at, not what it says. The Companion app is built, and on test phones only.

The Comm app's new AI area: "ANTON with friends"

The newest piece, and the smallest. In October 2026 we built the first step of an AI area inside the Kira communication app. It uses **your own key** with an AI service you choose, straight from the phone — not through the relay. It offers everyday expert packs derived from ANTON's modules (travel, food and home, learning, money basics), shows **"what went in"** for every answer — model, service, privacy setting, tokens, cost and the exact messages — and checks for personal data before anything is sent: Swedish personal identity numbers, card numbers, IBANs, e-mail addresses and phone numbers. A wallet recovery phrase cannot be sent at all. On the default service it asks for zero data retention and no data collection, it has a daily spending cap, and it is switched off for child profiles.

And it does nothing. **It never sends, books, pays or agrees.** It answers. Professional packs, the family's experts as "friends", and actions that a person confirms are planned next, in that order. This first step is on test phones.

Portals and Beehive

Two more pieces live inside ANTON itself. **Portals** are web spaces a user builds in ANTON that are both a page for people and an endpoint other software can call, with its capabilities declared. Building, publishing, visiting and calling one works end to end. **Beehive** lets several ANTONs deliberate on one question in rounds, with signed contributions, dissents preserved rather than smoothed away, and a synthesis exported as an .anton evidence bundle; a boundary decides which knowledge each participant may see. Beehive is built in ANTON's code; it has not yet been proven between separate organisations.

Status at a glance

- **In main code** (merged, not checked as deployed): Portals; Beehive (not yet proven between organisations); the relay, as a package in ANTON's public repository.
- **Built, not released:** the Kira app (Android, not in Google Play); the ANTON Agent app; the Collaboration and Agent Pay gateways (local install); ANTON-to-ANTON messages and mission tasks; mesh v0.1; the Companion app and the Comm app's AI area (test phones).
- **Designed or planned:** an independent external security review; a mesh enterprise profile; "borrow from a project"; the next steps of the Comm AI area; a shared library of .anton bundles.

Why a relay — and why you can read it

A fair question is why there is a relay at all. Why not let installations talk to each other directly?

The answer is practical. A direct connection means one side has to accept incoming connections from the internet — an open door on a machine that holds compliance work. Most firms, rightly, do not allow that, and most home and small-office networks make it hard anyway. With a relay, **both sides dial out**. Nothing has to listen for strangers. The relay simply meets the two outgoing connections in the middle and passes sealed envelopes between them.

That makes the relay the one piece every message passes through, so it deserves the most scrutiny. Three things help. First, it is built to carry only sealed envelopes: it can see the addresses and the times, not the contents. Second, it is a separate package inside ANTON's open-source repository, under the same licence, the EUPL, so anyone can read exactly what it does — and anyone can run their own. Third, it has its own tests for the failures that matter most in a shared carrier: tampered messages, one tenant's traffic leaking to another, someone squatting on another's address, messages sent to the wrong place, and attempts to get round its rate limits. The mesh, for its part, comes with a written threat model that sets out what it is designed to resist and what is out of scope for this version.

None of that is an independent review. It is the material an independent reviewer would start from — and it is open, so the review does not depend on our goodwill.

The family stays on files, on purpose

One thing may look like a contradiction. In *Not one big button* I wrote that the family's specialist experts — Valentin, Olivia, Horizon and the others — exchange work only as **files that a person approves and carries**, with no hidden live calls between them. Now I am describing live, encrypted links between ANTON installations. Which is it?

Both, deliberately. The specialist experts handle regulated deliverables — signed reviews, client annual cycles, change registers — where every hand-off must be visible and approved, and a few seconds of friction buys a clear record. That agreement stands, and nothing in this paper changes it. The relay and the mesh belong to ANTON, the generalist, and to the people and agents around

it. They are the right tool for asking a colleague's expert a question or handing over a task between consenting installations — with signatures, and with a person approving every commitment. **Different kinds of work, different kinds of connection, the same rule underneath: nothing important moves without a person seeing it.**

Four uses of a connected network, honestly

When we planned this series, we listed four things a connected network should let people do. Here is where each one stands.

Use	What exists today	Status
Asking a colleague's ANTON	Ask another installation's specialist agents; the answer is produced on that installation, with its own model and knowledge	Built; via a signed queue, roughly every 30 seconds
Borrowing knowledge from an expert's project	Rules decide what may go to peers: only knowledge marked as shared, never private project knowledge; a boundary selects what each peer may see	Rules built; a "borrow from a project" experience is planned
Agent-to-agent requests	Agents find each other, ask, negotiate and reach a signed agreement; a person approves each commitment	Built in the standalone gateways; purchases settled only in testing
Installation to installation	Two ANTONs exchange signed, encrypted messages and hand over tasks with a signed brief and result	Built

Two of these deserve a word.

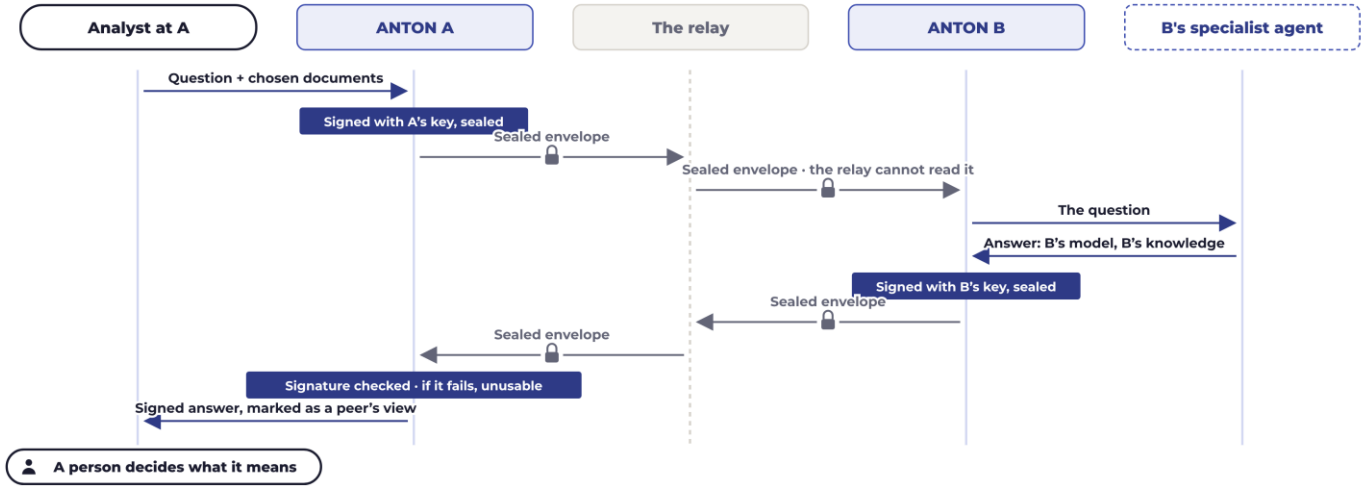
Borrowing knowledge. The idea is attractive: an expert's project has learned something, in the form of the small pieces of learning we called atoms in WP8, and a colleague could borrow it. The rules that make this safe exist — only what is explicitly marked as shared can ever go to a peer, and private project knowledge never does. But atoms themselves arrive slowly and are gated, as WP8 explained: ANTON adds nothing from its memory to a prompt until enough real work and ratings exist, and it is still collecting. So the honest status is "the rules are built; the experience is planned".

Agent-to-agent requests in finance. When we first sketched this paper, the examples were a customer investigation, a file, transactions — one institution's agent asking another's. Those are exactly the requests that touch bank secrecy, the AML rules on information sharing and data protection. They are **not built**, and they should not be built until there is a clear legal basis, the secrecy rules are respected and four eyes are in place on both sides. I include them here as a future possibility that would need all of that first — not as something on offer.

Part 4 · A worked example

Asking another ANTON's expert

Two fictional organisations: a small advisory firm, **A**, and a specialist in Nordic sanctions, **B**. Both run ANTON. They connected earlier by exchanging contact codes in person. An analyst at A needs a second view on whether a Finnish sanctions listing affects a client's supplier.



Signed out, signed back, read by nobody in between.

Step	What happens	Who acts	Status
1. The question	The analyst opens "Ask a peer's agent" and picks B's sanctions specialist agent from B's public directory.	Person	Built
2. What leaves A	Only the question she typed and the documents she chooses. No client name: she uses a fictional label. The message is signed with A's installation key, encrypted end to end once both sides hold keys, and carried by the relay as a sealed envelope.	Person, then code	Built
3. At B	B's agent answers on B's installation, with B's own model and knowledge. Nothing of B's private knowledge leaves except the answer.	AI at B	Built
4. Back to A	The signed answer arrives through the queue. A's ANTON checks the signature. If it fails, the answer is marked unusable.	Code	Built
5. A person decides	The analyst reads the answer as a second opinion, not as a fact, and records it with its source: "peer: B's sanctions agent, signed".	Person	Built

Step	What happens	Who acts	Status
6. Optional: a task	If A wants B to do a piece of work, A hands over a mission task with a signed brief. B's operator accepts or declines, and A is notified. The result comes back signed.	Person at each end	Built
7. The record	Both sides keep what was asked, by whom, when, and the signatures.	Code	Built

What the round trip shows

Go back to Scene 1. The analyst no longer has to wait three days for a summary from memory. She can ask the colleague's expert directly, and the answer is produced where the knowledge lives — with that installation's own model and sources — and sent back signed. The colleague's private work never leaves her workspace. Only the answer does.

Notice where control sits at each step. **A decided what left A.** The question and the chosen documents, and nothing else — not the client's name, not the project, not the rest of the file. **B decided what B's agent could draw on. The relay carried two sealed envelopes** and saw only that two installations exchanged messages at a certain time. And at the end, **a person decided what the answer meant.** Not the agent that wrote it, and not the system that delivered it.

Notice also what the record gives an auditor, compared with Scene 2: not a forwarded email chain, but a signed question, a signed answer, the time of each and the source — on both sides. Not a complete answer to every question about sharing between firms, but a far better starting point.

The relay carried it, neither side lost control, and a person decided what it meant.

Part 5 · From here

What this means for you on Monday

1. **Map where expertise sits** in your organisation: which people, teams and tools hold knowledge others need. That map is your network, before any software is involved.
2. **Write a sharing rule.** What may be shared between teams, what with partners, and what never. Make "shared" an explicit choice, never the default.
3. **Ask every agent vendor one question:** can the agent approve its own spending or its own agreements? If the answer is anything but a clear "no", stop.
4. **Ask for metadata honesty.** "End-to-end encrypted" tells you about content. Ask what the relay or the server can see — who, when, how much — and how long it keeps it.
5. **Wait for the review.** Treat any new communication channel as unapproved for client data until it has had an independent security review. Ours included.

Looking back: six gaps, eleven papers

We started in October with a simple claim: AI has not delivered in compliance because of the gaps around the model — **knowledge, time, training, trust, safety and governance** — not because of the model.

The first act was about how we think. In *No mystery box* we argued that an answer you cannot explain is worth nothing, and that the AI should find, a rule should rate and a person should decide. In *Not one big button* we argued for experts close to the work, with a generalist for the long tail. In *Trust the process*, that confidence comes from roles, checks and sign-off, not from chatting.

The second act was the method. *The seven layers* showed where the knowledge comes from and how it is stacked. *Setup beats size* argued that the right setup often matters more than the largest model, with honest limits on what we have measured. *Your knowledge, your side* argued that models are utilities, and your data and way of working must stay yours. *Atoms* described a learning loop in which each task can make the next one start better — carefully gated, and owned by you.

The third act was the new way of working. *From once a year to all the time* argued that checks can follow change rather than the calendar. *The .anton format* gave expertise a package that can travel and be checked. And this paper has asked how it can travel safely, between people, teams, agents and organisations.

Notice that not one of those papers argued for a bigger model. Every one of them was about what surrounds the model: method, knowledge, rules, records, people who sign. That was the thesis in October, and having written it out across eleven papers, I believe it more firmly now than when I started.

Where we believe this is going

Before the plans, I want to say where I think all of this leads. What follows is opinion — what I believe after writing these eleven papers, and after over fourteen years of doing the work they describe. It is the direction we are building in, not a description of anything that exists, and I would rather you judged it as an argument than read it as a promise. **Within a few years, I believe, good compliance work will look different in five ways.**

It will run continuously, and every conclusion will carry its own explanation. The annual clock was set by the cost of looking, and that cost is falling; when a check costs hours rather than weeks, the honest question is what should trigger it, not how often we can afford it. That was *From once a year to all the time*. But checks that follow change are only worth having if each result can be read and defended, so the second half matters as much as the first: what went in, which rule applied, who decided — not reconstructed when the auditor asks, but carried with the conclusion from the moment it is made. *No mystery box* and *Trust the process* argued for that standard. I expect supervisors and boards to take it for granted within a few years, and to be right to.

Organisations will own their professional context, and buy their models like utilities. By professional context I mean what makes your work yours: your method, your knowledge, your records, and the learning each finished task leaves behind. The models will keep improving and keep changing hands, and you will want to switch between them the way you switch electricity suppliers — caring that the power is reliable and fairly priced, not building your house around one company's plug. *The seven layers of knowledge*, *Setup beats size*, *Your knowledge, your side* and *Atoms* were four ways of making the same point: the value sits around the model, so that is what you must hold.

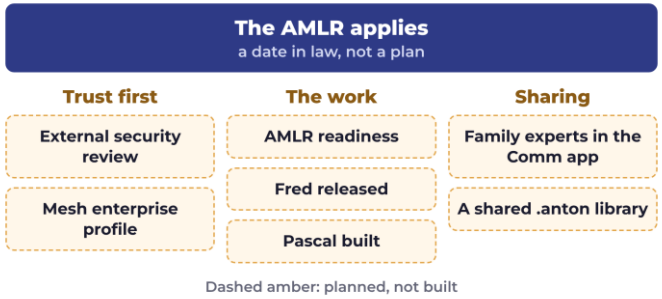
Expertise will travel. Packaged, checked and signed, between colleagues, teams and firms, so that a careful method built in one office does not have to be rebuilt in the next, and the knowledge that is almost always partly somewhere else can be asked for without being handed over. *The .anton format* and this paper described the first pieces. What I hope for is a network in which every node keeps its sovereignty — and in which Europe sets its own terms for how professional knowledge is protected, how agents may act and who answers when something goes wrong, rather than accepting terms written somewhere else.

And people will spend their time on judgement, improvement and the conversations that matter. This is the one I care about most, and it is where the series began: in *The six gaps* I called it a change of rhythm, not of headcount, and in *Not one big button* I argued for experts close to the work rather than one machine for everything. AI as a coworker that carries the routine — the compiling, the reconciling, the first draft — so that people can do what the routine always crowded out: deciding what matters, fixing the control instead of documenting its failure, talking to the product team before the launch rather than after it. A person still signs. That has never been the part I wanted to change.

> **Not a bigger model. A better context — owned by you, explained at every step, and decided by a person.**

That is what we mean by **artificial professional contextual intelligence** — APCI: intelligence that lives in the professional context rather than in the model, and that belongs to the people and organisations whose context it is. I cannot prove that this is where things are going; nobody can, and some of it will come faster than I expect, some slower and some differently. What I can do is say what we believe and why, build towards it one honest step at a time, and mark clearly, at every step, what is built and what is still a plan.

2027: what we plan



2027: plans, not promises. Trust first.

These are the next steps towards that. Each is a **plan**, not a commitment, and carries no date unless the date is the law's.

The AMLR applies from 10 July 2027. That is a fact, not a plan, and it shapes everything else. In the first half of 2027 the family's focus is helping teams get there: gap assessments, the BWRA, policies, validations and continuous follow-up.

Trust first. We intend to commission an independent external security review of the relay, the mesh and the agent gateways, and will make no claim of fitness for client work before it is done. After that, a mesh "enterprise profile" is planned — keys held in hardware, mutual authentication, signed audit-log export and single sign-on — for firms that must inspect their own traffic.

The work itself. Fred, the investigation wingman, is planned to move from its first slice to release. Pascal, the work board for people and AI experts together, is planned to be built from its design.

Sharing. The AI area in the Comm app is planned to grow: professional packs, the family's experts as "friends", and actions that a person confirms. And a shared library of .anton bundles between organisations, with ratings — and later, perhaps, paid expertise — is where last week's universe could start to grow.

In January

The series continues. In January we start with a paper on **AI governance you can audit** — how a compliance team using AI can show a supervisor what went in, who decided and who signed, against the EU AI Act, DORA, GDPR and the rules on outsourcing.

A closing thought

The model was never the problem; the gaps were, and we can close them ourselves. Connection is the last piece of that work, and the most human. The best compliance work has always depended on someone else — a colleague in another office, a reviewer, a counterpart at another firm. Done well, AI makes them easier to reach and leaves a record of what was asked and answered. Done badly, it puts everything in one place and lets software decide. We have tried to build the first kind, and to be honest about how far we have got.

Thank you for reading, for arguing with these papers and for telling us where AI stopped working in your team. I hope the year ahead brings you fewer gaps and more time for the work you never had time for. Merry Christmas, and see you in January.

— 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", and nothing described in this paper has yet had an independent external security review, so keep client work on your firm's approved channels until it has. 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.