

ANTON by openEXPERT

The ANTON family whitepapers · Autumn 2026 · Paper 3 of 11

Not One Big Button

Why real compliance work needs experts close to the people who use them — and how they work together

> *A generalist for the long tail. Specialists for the work that must stand up in front of a supervisor. One way of working for all of them.*

Version 1.0 · 29 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

Where we left off

In the first paper we named six gaps between a capable model and compliance work someone can stand behind. In the second we went deep on trust, and argued that the box has to be glass: what went in, how every number was computed, and who decided must all be open to anyone who asks. This paper turns to the knowledge gap — how the expertise gets into the tool, so that the person using it does not have to supply it every time.

This week we turn to a question every team asks sooner or later, usually in a procurement meeting: **should we buy one AI that does everything, or several that each do one thing well?**

It is a fair question, and the honest answer is not the one most vendors give. The answer is *both* — but in very different roles. And the reasons why say a great deal about how professional work actually gets done.

The dream of one button

Every AI pitch I have seen in the last two years ends in the same place. One box. Upload everything, ask anything, get the answer. One button for the whole compliance function.

I understand the appeal completely. Compliance teams are stretched thin. Tools multiply. Every new system means another login, another training session, another integration nobody has time for. The promise of *one* thing that just works is enormously attractive, and part of it is true: a good general AI really can help with an astonishing range of tasks.

But in a compliance function, the one-button dream breaks — and it breaks in predictable places. This paper is about where, why, and what we built instead.

In this paper

- **Part 1, the problem:** where one big button breaks, and three scenes you may recognise.
- **Part 2, our thesis:** a generalist and specialists, and why "opinionated" is a feature.
- **Part 3, our take:** five principles, the family, who to ask for what, and how the experts work together.
- **Part 4, a worked example:** a new EBA guideline lands, followed through five experts.
- **Part 5:** what this means for you on Monday, and next week.

Part 1 · The problem

Where one big button breaks

A general AI fails in a compliance function in three specific ways. None of them is about intelligence.

The defaults are wrong. A general tool starts from nowhere in particular. It does not know that a payment institution and a bank carry different obligations under the AMLR, that a fund company's annual compliance plan looks nothing like an insurer's, or that a model validation report needs a conclusion on data quality before it can say anything about thresholds. Every one of those things has to be explained, every time, by the person using it — and that person usually discovers what was missing only after reading the result.

The words are wrong. Compliance has a language, and the language carries meaning. The supervisor's guidance says *residual risk*; the general tool says *risk score*. The board *adopts* a policy; the tool says it was *approved*. A finding is *accepted* by a reviewer, not *confirmed by the AI*. These sound like small differences. To a supervisor, an auditor or an experienced MLRO, they are signals — of whether the author understood the framework or merely paraphrased it.

The flow is missing. This is the deepest one. A gap assessment, a model validation and a business-wide risk assessment are not questions with answers. They are *processes*, with steps, roles, evidence and sign-offs. Who prepares, who reviews, who decides, what is frozen when, what goes to the client for a factual check. A chat window has none of that. It has a question and an answer, and everything else is left to the user to organise around it.

> **The blank chat box is the most expensive screen in your compliance function. It offers everything and suggests nothing.**

Three ways the button fails

Scene 1: the blank box

A newly appointed compliance officer at a small fund company opens a general AI assistant to *"do our annual compliance plan"*. The blank box offers everything and suggests nothing. Which regulations apply? Which cadence? Who owns what? What does the supervisor expect to see in March?

After an hour of prompting, she has a long, generic plan that reads like every other annual plan ever written. It is not wrong, exactly. It is just not *hers*.

What she needed was a tool that starts from her type of firm, gives her a dated annual cycle with owners and evidence already in place — and then lets her adjust it with her own judgement.

Scene 2: the "everything" platform

A risk team buys an enterprise AI platform with hundreds of features. Six months later, the team uses two of them. Nobody is quite sure which settings matter. Every report comes out looking slightly different, depending on who ran it. When the team raises this, the vendor's answer is more training.

What they needed was fewer choices — with the right defaults — for the three jobs they actually do every month.

Scene 3: the hand-off that never happened

The regulatory change team spots a new EBA guideline. The policy owner hears about it three weeks later, in a meeting, almost by accident. The reviewer who signs off the policy update never sees the guideline at all.

Three people, three tools, no thread between them. Each did their part well. The work still fell through the gaps between them.

What they needed was experts that pass work to each other in a known format, with a person approving every hand-off.

The paradox of the powerful tool

There is something counter-intuitive in Scene 2 that is worth sitting with. The platform was not weak; it was *too* capable. Hundreds of features, every option configurable, nothing assumed.

Power without defaults is a burden. Every choice a tool leaves open is a choice the user has to make — and in a busy compliance team, most of those choices get made inconsistently, by different people, on different days. The result is variation, and variation is exactly what a regulated function cannot afford.

The most helpful thing a professional tool can do is make most of the decisions for you — the right ones, visibly, so you can change them when you need to. That is not a limitation. It is expertise, built into software.

Part 2 · Our thesis

> **One generalist, many specialists, one way of working.**

The thesis

A single "do everything" AI is a good start and a poor finish. For work that must be defensible, an expert that sits close to the work wins, because it brings five things a generalist cannot:

1. **Fewer choices.** It asks only what matters for this job.
2. **The right defaults.** The first run is already structured for this kind of firm and this kind of reader.
3. **The right words.** It speaks the supervisor's and the practitioner's language, in English and Swedish.
4. **A guided flow.** The steps, the roles and the sign-off are built in, not left to the user.
5. **Its own rules and record.** It knows which numbers must come from rules, and who must confirm them.

But — and this matters just as much — the generalist does not go away. There is a long tail of work in every compliance function that will never justify a dedicated tool: the one-off memo, the quick comparison, the first draft of something new. That is where a generalist earns its keep. The mistake is not using a generalist. **The mistake is using a generalist for the work that has to stand in front of a supervisor.**

Why we believe it

We built both, and watched what happened. ANTON can do almost everything — it has more than 550 expert modules across 60 areas, and it is genuinely good at most of them. And yet, every time a job needed to stand in front of a supervisor — a gap assessment, a review before signing, a business-wide risk assessment — we found ourselves building a focused expert for it. Not because ANTON could not produce the document, but because the *work around the document* needed its own flow, its own rules and its own record. Each focused expert ended up simpler for the user and stricter underneath. That pattern repeated too many times to be a coincidence.

Practitioners work in roles, not features. Nobody in a compliance function wakes up thinking about features. They wake up as an MLRO with a BWRA due, a model validator with a report to sign, an internal auditor with a review to plan, an outsourced desk with twelve clients' annual cycles to keep on track. When we built the family portal, we started its "Find your expert" guide from the visitor's role, not from a product list — and almost every role turned out to need two or three experts, never all of them.

Opinionated is a feature. This is the reason I care about most. Each of our focused experts has opinions — things it simply refuses to do:

- **Gabriel** will not let the AI decide a score. A fixed rubric does that.
- **Valentin** will not review a document with the same model that wrote it.
- **Olivia** will not put two clients' information in one AI request.

These are not limitations to apologise for. **They are the guardrails that make the output safe to rely on.** A tool with no opinions about what it refuses to do is a tool with no guardrails — and in regulated work, that is not flexibility. It is risk.

An analogy from medicine

Think about how healthcare is organised. A good general practitioner can help with a remarkable range of problems, and most of what most people need is handled there, well. But when the question is a heart valve, you see a cardiologist — not because the GP is not clever, but because the cardiologist has the defaults, the language, the procedures and the accountability built around one kind of problem.

Nobody thinks the existence of cardiologists makes GPs obsolete. Nobody thinks a hospital would be better if every specialist were replaced by one very capable generalist with a long checklist. The system works because both exist, in the right roles — and because there is a known way to refer a patient from one to the other, with the notes going along.

That last part — the referral, with the notes — is the part most "AI platforms" miss. And it is where we will end this paper.

What this is not

It is not a collection of disconnected tools. That would just be Scene 3 with better software.

The family shares **one way of working** — AI proposes, code computes, a person decides — **one contract** for "what went in", **one set of platform rules**, and **known hand-off files** between products. Different experts, one discipline.

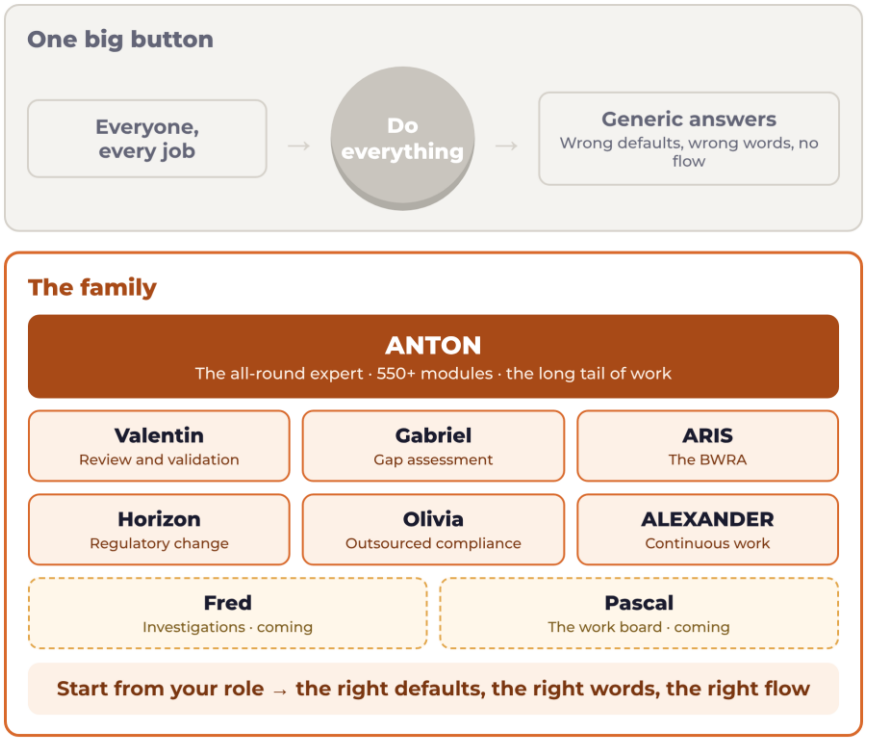
Part 3 · Our take: the family

Five principles

1. **One generalist, many specialists.** ANTON for breadth; a focused expert for each job that must be defensible.
2. **Close to the user.** Each expert is built for one group of people and their daily work, in their words.
3. **Simple on top, strict underneath.** A short, guided flow on the screen; rules, four eyes and a trail in the database.
4. **Hand-offs, not hidden links.** Work moves between experts as a file a person approves and carries. There are no hidden live calls between products.
5. **One way of working.** The same motto, the same "what went in", the same rule that the AI never sets a score.

Principle four is the one people question most, so let me explain it. It would be technically easy to wire the products together so that work flows silently from one to the next. We chose not to. **Every time work crosses a boundary between experts, a person sees it, approves it and carries it.** That is slower by a few seconds. It is also the reason that, months later, anyone can see exactly what moved, when, and who let it through. In regulated work, an invisible hand-off is an unaccountable one.

Meet the family



One generalist, many specialists, one way of working.

Picture: The family's character portraits, in a row (from the portal)

ANTON — the all-round expert. For anyone in compliance, risk, legal or finance. More than 550 modules in 60 areas; drafts, analyses, ANTON's own Risk Atlas for any business, and exports to Word, Excel, PDF and PowerPoint. ANTON can still do almost everything, and it is open source.

Valentin — review and validation. For the second line, internal audit, model validators and consultants. Independent review and a second opinion before you sign; audits; model validation from thresholds and metrics through to a signed report and a board summary.

Gabriel — the Gap Assessor. For compliance, AML, risk and audit. Article-by-article gap assessment against 62 frameworks. A fixed rubric scores, and a different model gives a second opinion.

ARIS — the business-wide risk assessment. For MLROs, FCP analysts, approvers and audit in banks. The BWRA under AMLR Article 10, with every figure traceable to a person, a rationale and a method. (ARIS is a separate product from ANTON's Risk Atlas: ARIS is built for banks' BWRA; the Risk Atlas serves any business.)

Horizon — regulatory change. For FCP, compliance and risk teams. Regulatory horizon scanning for financial crime: the law's own words, deadlines, lessons from enforcement, and a change register with four eyes.

Olivia — outsourced compliance. For outsourced and co-sourced compliance teams. One walled-off workspace per client; each client's *årshjul* (annual compliance cycle) built from baseline templates; four eyes; and a view across all clients at once.

ALEXANDER — continuous work. For compliance, risk, AML and audit on a standing mandate. Keeps a mandate under review, re-checks on a schedule or when something changes, and shows what moved since last time.

All seven are live on public showcases with fictional data, gathered at ai-portal.futurechain.eu. Only ANTON is open source.

Coming, and marked as such

- **Fred — investigations.** The investigation wingman for financial crime investigators: works a transaction-monitoring alert from start to finish, with code computing the likelihood band. A first slice is built and works on synthetic data with the AI switched off. It is not public yet.
- **Pascal — the work board.** A board for people and AI experts together: who did each task (a person, an AI expert, or both), estimates against actuals, and what was learned. Being designed.

Who to ask for what

Start from your role, not from a feature list. This is the logic of the portal's "Find your expert" guide, in short:

If you are…	Ask first	Then
Compliance officer	ANTON — policies, memos, plans	Gabriel for gaps; ALEXANDER to keep it under review
MLRO	ARIS — the BWRA	Horizon — what changes, and when
Model validator	Valentin — validation to a signed report	ARIS — the risks the models must cover
Internal auditor	Valentin — audits and independent reviews	Gabriel, ARIS
Outsourced compliance desk	Olivia — each client's annual cycle	Valentin before it goes out; Horizon
Reviewer before sign-off	Valentin — a second opinion bound to the text	Gabriel — a second model's view
Regulatory change	Horizon — the law's words and deadlines	ALEXANDER — what it means for a mandate
DPO	Gabriel — GDPR, article by article	ALEXANDER, ANTON
Cyber and operational resilience	Gabriel — DORA gap assessment	ALEXANDER, Horizon
AML investigator	ANTON today; Fred when released	Horizon

Notice that nobody's row has more than three names in it. That is the point.

How the family works together

A family of specialists is only better than one big button if the specialists can work as one. Here is how that works today.

One way of working. Every expert follows *AI proposes, code computes, a person decides*, records "what went in" in its own form, and enforces four eyes.

One written platform agreement. Valentin, Olivia and Horizon share a written agreement on how they behave, held word for word in all three codebases — with a test in each that fails if a copy drifts. Shared building blocks have one "home"; copies are made in one direction only and carry a header naming where they came from.

Hand-off files a person approves. Every exchange between these products is a file that a person approves, carries, and the receiving side reviews. Olivia can send a request for a review to Valentin; Valentin can start the review from that request and send back a signed receipt, which Olivia checks against its fingerprint before filing it as evidence. Horizon publishes change items in a shared format for the others to use. There are no live calls between the products.

Shared method, not shared confusion. Olivia runs a set of ANTON's modules copied word for word at a fixed version, with a check that warns if they drift. And Olivia's "next step" suggests, by rule, which module or which sibling expert fits a gap it has found.

Coming: in its design, Pascal takes work items from the others — Horizon's changes, Olivia's annual duties, Valentin's findings, Gabriel's gaps — through one versioned contract, and sends status back.

Status at a glance

- **Live** (public showcases, fictional data): ANTON, Valentin, Gabriel, ARIS, Horizon, Olivia and ALEXANDER; Horizon's feed and change register; Olivia's annual-cycle items; ALEXANDER's triggers.
- **Built, not released:** Fred's first slice (synthetic data, AI switched off); the Olivia–Valentin review request and signed receipt.
- **Designed or planned:** Pascal and its work-item contract; the Horizon-to-Olivia hand-off (format agreed, parts still being built).

Part 4 · A worked example

A new EBA guideline lands

The example is fictional: *Exempelbanken* (Horizon's demonstration bank) and an outsourced compliance desk that serves it. It shows how the family is **designed** to handle one regulatory change from start to finish. Some hand-offs are built today and some are still being built, so each step is marked.



Five experts, one thread, a person at every hand-off.

Step	What happens	Expert	Status
1. Spotted	The guideline appears in Horizon's feed from the official source, labelled with its authority and legal force, and linked to the AMLR duties it touches. A person accepts the reading; a second person approves the change-register entry.	Horizon	Live
2. Handed on	Horizon exports the change as a file: the change, its date, a summary, a link and a fingerprint. A person carries it on.	Horizon → Olivia	Format agreed; parts built
3. A duty opens	The desk imports the change and adds a dated item to the client's annual cycle, with an owner and the evidence needed.	Olivia	Cycle items live; import from Horizon being built
4. Drafted	The policy owner drafts the update with an ANTON module, with the guideline's knowledge pack and the house style as layers.	ANTON	Live
5. Reviewed and signed	Olivia sends a hand-off request. Valentin reviews the draft with a model that did not write it, checks the quotes, and a reviewer who is not the preparer signs. Valentin returns a signed receipt.	Olivia → Valentin → Olivia	Built
6. Filed as evidence	Olivia checks the receipt's fingerprint, and marks the duty done only when a person confirms.	Olivia	Built
7. Kept under watch	On a standing mandate, the new policy version is a trigger for ALEXANDER to re-check what moved since last time.	ALEXANDER	Live (triggers)
8. On the board	Who did each step — a person, an AI expert or both — and how long it took against the estimate.	Pascal	Planned

What the thread changes

Go back to Scene 3, where three people did their part well and the work still fell through the gaps. In this version, nothing depends on anyone happening to mention the guideline in a meeting. The change arrives as a file, becomes a dated duty with an owner, is drafted with the right knowledge, is reviewed by someone who did not write it, and comes back as signed evidence — and a person approved every step where work crossed from one expert to another.

That is the real argument against one big button. It is not that a single tool cannot write a policy update; it probably can. It is that compliance work is a *chain of responsibility*, and a chain needs links that can be seen. **Five experts, one thread, and a person at every hand-off** is not more complicated than one button. It is simply honest about how the work is actually done.

Part 5 · From here

What this means for you on Monday

1. **List your three most important recurring jobs** — the annual BWRA, policy reviews, regulatory change, say. For each, ask: does our AI tool know the defaults, the words and the flow of this job? If not, that job needs more than a chat box.
2. **Map roles to tools, not features to people.** Start from who does the work, and give each role the one or two tools that fit it.
3. **Write down your hand-offs.** Where does work pass from one person or team to another? Make each hand-off a named document with an approver. That is where things get lost today.
4. **Keep one generalist.** A general expert is still the right place for the long tail of one-off tasks. Just do not use it for the jobs that must stand in front of a supervisor without the right flow around them.
5. **Ask vendors for their opinions.** What does the tool refuse to do? Which numbers will it never let the AI set? A tool with no answer has no guardrails.

Next week: trust the process, trust the output

Organisations do not lack trust in AI because the answers are bad. They lack it because the work happens in ad hoc chats. Next week we look at how roles, evidence, checks, the client's factual-accuracy round (*sakgranskning*) and a signed, frozen report turn AI help into something an organisation can rely on — and how that maps to the EU AI Act, DORA and GDPR.

A closing thought

The one-button dream is really a dream about simplicity, and simplicity is a good thing to want. But in professional work, simplicity does not come from one tool that does everything. It comes from each tool doing one thing so well, with such good defaults, that using it feels effortless — and from those tools knowing how to pass the work along.

We did not set out to build a family. We set out to build one very good generalist, and the work kept telling us that some jobs deserved their own expert. I think that is the most honest way to arrive at an architecture: not from a diagram, but from listening to where the work actually got stuck.

— 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; 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.