

AI in compliance: what already works today.

32 compliance leads showed us where their time disappears. This paper sets out which of these tasks AI models can take on in 2026, which they can't, and what the first step looks like.



CONTENTS

What this paper covers.

Part I covers the survey in six pages. Part II translates it into four tasks that absorb working time. Part III shows which of these tasks today's AI models can take on, and which they can't. Part IV describes the path to a first implementation.

PART I · WHAT THE SURVEY SHOWS

1 You no longer need a software project	03
2 The situation, in numbers	04
3 Our sample at a glance	06

PART II · WHERE THE TIME GOES

4 The biggest lever is reporting	07
5 Your data is PDFs, not a database	09
6 The wish for in-house collides with the market	10

PART III · WHAT AI MODELS CAN SOLVE

7 Chat, a folder with a workflow, a bespoke application	11
8 Your tool can already do this	12
9 Case study A: the certification navigator	14
10 When a bespoke application pays off	16
11 Case study B: the sustainability one-pager	17
12 What matters is where your documents sit	19
13 Risks that limit deployment	21

PART IV · WHAT GETTING STARTED LOOKS LIKE

14 Optimise first, then digitalise	22
15 The governance minimum before the first workflow	24

APPENDIX

A Topline by question · B Sources · C Coding	25
--	----

03

01 · MANAGEMENT SUMMARY

You no longer need a software project.

First. Much of compliance work today is information logistics done by hand. **81 per cent** track regulatory change through manual research, **6 per cent** with a specialised tool.

Second. The effort involved is substantial. **One organisation in two** spends 20 hours a week or more on manual compliance tasks. Only 13 per cent stay under five hours.


Third. **56 per cent** don't measure how well their programme works. The reasons are technical, not cultural: 83 per cent of this group cite a missing data foundation, 61 per cent the lack of a suitable tool.

Fourth. Nobody is asking for another dashboard. **75 per cent** want to process data in-house or in a hybrid model, **84 per cent** work with unstructured documents. What's needed is a controlled process built on their own data.

Fifth. As of 2026, that process no longer needs a software project: a defined folder, a workflow held in a file, a traceable result. We show two implementations, and the limits that aren't up for negotiation.

02 · AT A GLANCE

The situation, in numbers.



The effort at stake




Figure 1: Weekly time spent on manual compliance tasks (n=32). Orange marks the 50 per cent at 20 hours or more, that is, 16 organisations. At n=32, one person equals roughly 3 percentage points, which is why we give the absolute number throughout.

One organisation in two spends half a working week on tasks that call for barely any judgement: gathering information, reformatting it, and passing it on.

How regulatory change is tracked today




Figure 2: Multiple answers allowed, n=32. 94 per cent currently work without a regulatory-intelligence tool.

What data is actually available



Figure 3: Multiple answers allowed, n=32. The raw material is text, and that decides which class of tool fits.

The gap between ambition and evidence



The 44 per cent who do measure mostly track operational figures: training completion rates, case numbers, contracts reviewed. Few measure in a way that could actually inform a management decision. That gap is precisely what limits buy-in at board level.

Why measurement doesn't happen (base: the 18)




Figure 4: Two of the three most common reasons are tool and data problems, meaning they're solvable.

03 · SAMPLE

Our sample at a glance.


This paper is based on an online survey of 32 compliance leads with ties to the Swiss and Liechtenstein markets, fielded between 18 March and 31 May 2026.

32 is a solid number for a questionnaire of this length. It covered 35 fields, including open questions on pain points, areas of law, and desired features. Filling it in takes about half an hour. That's the real strength of this data: all 32 submissions passed server-side validation, and on average only three of the 35 fields were left blank, mostly ones that didn't apply given an earlier answer. These aren't panel responses clicked through in a hurry; they're answers from people who took the time. The free-text answers show it: nobody skimming a survey writes 'too many false positive hits on our compliance filters, which then have to be cleared manually'.

We're strict about how we read the data. For multiple-choice questions we report both the number of mentions and the share of respondents, not a distribution, so the percentages regularly add up to more than 100 per cent. For conditional questions we state the reduced base explicitly, for example 'base: the 18 who don't measure'. Free-text answers are coded by theme; the scheme is set out in Appendix C.

Two points deserve to be stated plainly. The sample is a targeted selection of practitioners, not a random sample of the wider population. And it skews towards large organisations: 72 per cent have 300 or more employees. Findings on resources and governance maturity would likely look worse for SMEs, not better.

Sector profile



Organisation size



Figure 5: Sample profile (n=32). Financial services and industry each account for 28 per cent. Bear the large-organisation skew in mind when applying these findings to SMEs.

04 · THE NEED

The biggest lever is reporting.

Reduce the 35 survey fields to their common core and four tasks remain. They differ structurally, and only two can be reliably automated today. Part III builds on that distinction.

TASK 1

Knowing what has changed

Tracking regulatory change, filtering it, relating it to your own organisation. 81 per cent do this through manual research, 63 per cent through trade associations, 50 per cent through external advisers. 6 per cent use a specialised tool. From the free text: 'flooded with information and rules', 'keeping an overview of current requirements, including in other jurisdictions'.

Automatable? Partly. As a flagging function with a sourcing duty, never as legal advice. See Chapter 13.

TASK 2

Producing reports and evidence

The single most common need. 91 per cent name reporting as a core function, 59 per cent see reporting as the activity most likely to be automated. 72 per cent name the annual compliance risk assessment as an automatable process, the highest score of all thirteen processes asked about.

Automatable? Yes, readily. Recurring, clearly bounded, and the result is instantly checkable. The best place to start.

TASK 3

Being able to prove effectiveness

56 per cent don't measure. Within that group, 83 per cent cite a missing data foundation and 61 per cent the lack of a tool, both technical problems. Notably, 9 of the 16 organisations spending 20 hours a week or more still don't measure. A great deal of work happens, and little of it is evidenced.

Automatable? Yes, once Task 2 is running. Metrics fall out as a by-product of a structured reporting process.

TASK 4

Governing your own AI

For 94 per cent, AI regulation is on the agenda. 38 per cent have made initial enquiries, and 28 per cent each are either monitoring or have governance in place. The biggest hurdle sits with people, not technology: 53 per cent cite a lack of in-house AI governance expertise, and 47 per cent each cite classifying their own systems and unclear regulatory requirements.

Automatable? No. This is a matter of competence and decision-making. Tools only help once that's in place.

The most striking figure in the survey is a combination: six organisations have established AI governance and, in the same breath, name a lack of in-house expertise as a hurdle. Governance on paper is not yet competence on the ground.

Which processes people actually want automated

Question 11, multiple answers allowed, n=32. The median respondent names five processes. The need is broad, not narrow. The top four entries all belong to Task 1 and Task 2.





Figure 6: Automation requests by process. Orange marks the three processes that combine high manual load with clearly bounded, repeatable steps, making them the natural place to start.

What a cockpit needs to do



What should trigger a warning



Figure 7: The most-requested trigger is, of all things, the hardest one to build: AI-detected anomalies. Deadlines and specific matters are deterministic, and so the more honest place to start.

05 · THE DATA SITUATION

Your data is PDFs, not a database.

One number decides which class of tool fits: 27 of 32 name unstructured documents as their most important data source, meaning PDF contracts, Word policies, directives and minutes. Emails and manually maintained spreadsheets follow with 22 mentions each. Structured databases trail with 18.

22 of 32 combine both: manual research **and** unstructured documents as the main source. That's what the working reality looks like: people read text and write summaries.

Classic GRC and business-intelligence tools are built around the table. They expect fields,

thresholds and interfaces. That is precisely not where the need lies. A tool that reads text, cites its sources, and only then produces a table fits the data situation. One that presupposes a well-maintained database just pushes the work upstream.

That also explains why 15 of 32 want a fully automatic, real-time API connection, while 5 answer 'I don't know' at the same time: the appetite for integration is there, the technical precondition for it often isn't. Starting with the documents needs no interface at all.

The three systems that would need connecting

20/32

ERP system

19/32

CRM system

19/32

Microsoft 365 (SharePoint, Teams)

Only 10 of 32 name a GRC module, and only 8 name dedicated legal-tech tools. The relevant data mostly sits in systems organisations already have, above all in the Microsoft environment. That's good news: the first useful step needs no new platform, just access to a defined folder.

< > 24 AI in Compliance

Grid icons

List icon

Table icon

Fullscreen icon

More options icon

Share icon

Search icon

Suchen

Name	Änderungsdatum	Größe	Art
> Whitepaper	Heute, 15:53	--	Ordner
Whitepaper_KI_in_Compliance_MMIND_Bratschi_Entwurf_v2.pdf	Heute, 15:37	2.9 MB	PDF-Dokument
KI_in_Compliance_MMIND_2026-08-20.pptx	Gestern, 18:57	1.7 MB	Office...entation
B_Corp_Eligibility_Treuhand_Vaduz.md	Gestern, 08:29	15 KB	Markdown
> Beispiel B Corp	Gestern, 03:10	--	Ordner
generator_praesentation.js	Vorgestern, 16:25	29 KB	JavaSc...t script
KI_in_Compliance_MMIND_2026-08-20.pdf	Vorgestern, 16:25	180 KB	PDF-Dokument
certification-navigator.zip	Vorgestern, 16:25	5 KB	ZIP-Archiv
ai-bcorp-team-repo-update.zip	Vorgestern, 16:25	90 KB	ZIP-Archiv
Pruefbericht_und_Korrekturen.md	Vorgestern, 16:25	11 KB	Markdown
> Beispiel Konrad	Vorgestern, 14:59	--	Ordner
AI Compliance Webinar 2026 final.pptx	14.08.2026, 13:00	4.5 MB	Office...entation
European Sovereign AI Compliance Hub...EIC Accelerator Partner Discussion.pptx	08.07.2026, 08:50	105 KB	Office...entation
European Sovereign AI Compliance Hub - UK Partner Discussion.pptx	30.06.2026, 02:05	70 KB	Office...entation
comliance survey_teilnehmer-auswertung.md	04.06.2026, 11:51	121 KB	Markdown

Figure 8: A dedicated data folder set up for a compliance process. The starting point is a folder structure, not a system.

O6 · OPERATING MODEL


The wish for in-house collides with the market.

The clearest preference in the whole survey. 50 per cent want AI processed exclusively in-house, a further 25 per cent want a hybrid model, 75 per cent combined. Only 19 per cent consider an external cloud solution acceptable. On hosting, 11 name a private cloud and 7 name on-premises; only 5 name SaaS. 13 want the operation run by their own IT.

This finding collides with the market. Nearly every capable AI tool is a cloud service. Anyone who takes 'exclusively in-house' literally ends up self-hosting open model weights, and accepts a measurable performance gap in exchange.

For data under professional secrecy, there's a further point many overlook: under the guidance note issued by the Zurich cantonal data protection commissioner in February 2026, US authorities can access personal data under the CLOUD Act even where it sits in data centres in Switzerland or the EU/EFTA area. **Swiss data residency with a US provider does not solve this problem.**

Where processing happens



Access model




Figure 9: 66 per cent require granular access rights. Combined with the preference for in-house processing, that adds up to a requirements profile no standard SaaS subscription meets.

What looks contradictory here is resolvable, but only if the question is reframed. Not 'which platform should we buy?' but 'which one process, with which data, do we bring under control?'

A word on shadow AI

The 'exclusively in-house' finding describes a wish, not the practice. Where an organisation fails to govern access to suitable tools, staff turn to freely available services, company data and all. According to IBM's 'Cost of a Data Breach 2025' report, 63 per cent of 600 organisations surveyed had no AI governance policy; 20 per cent had already recorded an incident linked to shadow AI, at an average extra cost of around US\$670,000 per case. A ban with no alternative shifts the risk. It does not reduce it.

07 · LEVELS OF SOLUTION

Chat, a folder with a workflow, a bespoke application.

When people talk about AI in compliance today, they usually mean Level 1: a chat window. The need identified in Part II sits at Level 2. Level 3 stays the exception, and Chapter 13 explains why.

LEVEL 1

Assistance: the chat window

One question, one answer. Useful for drafting, translation, a first read on something. **It doesn't solve the tasks from Part II:** nothing is saved, nothing is repeatable, nothing is checkable. Run the annual risk assessment through a chat and you start from scratch again next year, with nothing to show anyone about how the result came about. From the standpoint of Task 3, Level 1 is worthless: it produces no evidence.

LEVEL 2 · THIS IS WHERE THE LEVER IS

A folder plus a workflow held in a file

A defined folder holding the relevant documents, a workflow written down, a repeatable result. The workflow is a text file, not software. It sets out which sources to check, in what order, against which criteria, in what format the result takes shape, what must be cited, and where a human has to decide.

This is the answer to the survey. Repeatable, checkable, versionable, tied to a folder rather than a person. And buildable in days, not quarters. Chapters 8 and 9 show it in practice.

LEVEL 3

A bespoke application

A small purpose-built application with an interface, for recurring tasks involving many people or a large volume of work. Buildable today without a development team. **But:** from here on, operator and sometimes provider obligations arise, along with security responsibility and maintenance load. Chapters 10, 11 and 13 weigh up the benefit against the price.

Figure 10: The three levels don't differ in model quality. They differ in whether the result is repeatable and evidenced. Level 2 is the only place where both come together without a software product being built.

Repeatable?

Level 1 no
Level 2 **yes**
Level 3 **yes**

Evidenced for an audit?

Level 1 no
Level 2 **yes, via the folder and version history**
Level 3 **only with a purpose-built log**

Time to value

Level 1 minutes
Level 2 **days**
Level 3 weeks to months

08 · LEVEL 2

Your tool can already do this.

Four capability classes are enough to cover Task 2 and Task 3. All four are available across several tools in 2026, so the approach isn't tied to one provider. We name the products anyway, because otherwise nobody can order what they need.

CAPABILITY CLASS	WHAT IT DOES	WHERE IT'S AVAILABLE TODAY
Workflow held in a file a 'skill'	A text file describes a workflow, complete with criteria, format, a citation requirement and stop conditions. It loads automatically for a matching task and works through it. Versionable like any document.	Claude Skills · Codex and Gemini via AGENTS.md · GitHub Copilot via Custom Instructions · Mistral via Agents and libraries
Working in the folder	Reading, comparing and updating entire sets of documents where they already live, instead of uploading them file by file. Decisive for the 84 per cent working with unstructured documents.	Claude Cowork · Codex CLI · Mistral Vibe CLI · Grok Build CLI · Cursor · Gemini Antigravity
Persistent context a 'project'	A dedicated space with its own files, its own rules and its own history. The organisation's policy position lives there, not in one person's memory.	Claude Projects · ChatGPT Projects · Mistral Le Chat libraries · Copilot Spaces
Scheduled task	The workflow runs on its own: weekly, monthly, quarterly. A reminder becomes a process. A precondition for regulatory monitoring as business as usual.	Claude Tasks and the Cowork schedule · Codex Cloud in CI · any CLI in headless mode via cron

Figure 11: Capability classes rather than product names. Thinking in these four classes keeps you able to act when providers change. As at August 2026; product names and feature sets change fast.

What a workflow held in a file actually looks like

Excerpt, shortened

- Name:** regulatory-monitor
Purpose: Monthly regulatory report for the compliance function.
- Sources:** only the official pages and association notices listed in the source register. No others.
- Workflow:**
- 1 Read the source register.
 - 2 For each source, record changes since the last run.
 - 3 Check relevance against the organisation's topic list.
 - 4 For each hit: reference, date, verbatim quote, affected policy.
 - 5 Flag unclear cases as open questions rather than judging them.
- Limits:** no legal advice. Every assessment without a cited source is marked as such. Human sign-off is required before sending.

That's the whole trick, and it's precisely the point. A compliance workflow isn't programmed, it's **written down**. Anyone who can draft a policy can draft a workflow like this.

Three properties make it auditable, and all three are deliberate. First, **source discipline**: the set of permitted sources is closed. Second, **the citation requirement**: every statement carries a reference, or it's flagged as an assessment. Third, **the stopping point**: a human decides at a defined step.

Without these three properties, what emerges isn't a compliance tool but a text generator with compliance vocabulary. Chapter 13 puts numbers behind that difference.

What this looks like in practice

Three views from the certification navigator described in Chapter 9: the call, the workflow definition, and the repository through which an approved workflow is distributed across a team.




Figure 12: The call. The practitioner describes what they need in plain language. The matching workflow loads automatically. Nobody has to learn a command.




Figure 13: The workflow itself: a text file with a purpose, trigger phrases, steps, reference material and limits. Readable and editable by anyone who can draft a policy.

Why the repository is the lever

A workflow is a text file. That means everything true of any document under version control is true of it: changes are traceable, earlier versions stay retrievable, and approvals go through a documented review step.




Figure 14: The same workflows in a Git repository.

In practice, that means a workflow that group compliance has reviewed and approved once can be rolled out to every subsidiary through a repository. Change the policy, and the workflow changes in one place. That addresses the 66 per cent who require granular access rights, and the 22 per cent with multi-tenant requirements.

For an audit, that's the point that matters. Not 'we used AI', but: here is the approved version, here is the change history, here is the run that produced this report.

The difference from Level 1: what's left at the end isn't a piece of text, it's a record. Ask the same question again tomorrow and it follows the same path, checkable against the approved version.

09 · CASE STUDY A · LEVEL 2

A book of standards instead of three days of reading.

12

Employees, a PGR establishment, Vaduz

7.9_{MB}

Standards document as PDF (B Lab Standards)

1_{folder}

plus one workflow held in a file, no software

0_{lines}

of program code in use

The task

A fiduciary firm in Vaduz offering classic trust administration, bookkeeping, tax and asset management wanted to know whether, and under what conditions, B Corp certification was achievable. The basis is the B Lab Standards, version 2.2 of 20 February 2026: a several-hundred-page, versioned body of standards with Foundation Requirements, sector- and size-dependent scope, and deadlines in Year 0, 3 and 5.

Done by hand, that's several days of reading, with the real risk of missing a sector-specific extra requirement. It's exactly the process 22 of 32 respondents describe: read the text, judge relevance, summarise.

The set-up

A folder holding the standards document and a text version prepared for processing. Alongside it, a workflow held in a file, with a requirements matrix as a reference file, deliberately built so the matrix is only read once a specific requirement ID is actually needed.

The result

Eligibility report, as at 20 August 2026. Four of the five conditions in Foundation Requirement FR1 shown as met, one flagged as a decision for the board.

The real finding. B Lab has no 'trust and fiduciary' sector. But asset management falls under **investment advising**. That brings in two sector-specific extra requirements from Year 0: a documented human-rights review and an environmental review of investments, each covering the three most material holdings per year. Both apply regardless of company size. A 12-employee establishment faces the same requirement as a large asset manager.

And an honest gap. The report itself notes that the attribution logic changed between version 2.1 and 2.2, and that this point still needs verifying. That's not a flaw, it's the function at work: the workflow is instructed to flag uncertainty rather than paper over it.

What carries over elsewhere

This case isn't really about B Corp. It's the prototype for any task shaped like '**a public, versioned body of standards meets a specific company profile**', and so for a large share of the 72 per cent who name the annual risk assessment. ISO 37301, a FINMA circular, a parent company's code of conduct, a supplier requirement: same pattern, same approach. It works because of the structure of the source: public, versioned, tabular. That keeps every error checkable.

From folder to evidence




Figure 15: The stopping point at Step 4 is part of the design, not a concession. A workflow that knows its own limits is the only kind you can deploy in a regulated environment.

B Corp Eligibility Report

Unternehmen: Treuhandgesellschaft, Vaduz (Liechtenstein) **Rechtsform:** Anstalt (PGR) **Mitarbeitende:** 12 **Leistungen:** Klassische Treuhand · Buchhaltung und Steuern · Vermögensverwaltung **Operativ seit:** über 5 Jahre **Stand:** 20.08.2026 · B Lab Standards V2.2 (20.02.2026)

1. Ergebnis: Eignung bestätigt

Ihr könnt B Corp werden. Es gibt kein strukturelles Hindernis. Alle vier Bedingungen des Foundation Requirements FR1 sind erfüllt oder erfüllbar:

Kriterium	Status	Bemerkung
Mindestens 12 Monate operativ (FR1.1)	erfüllt	über 5 Jahre
Umsatz aus kompetitivem Markt (FR1.1)	erfüllt	Treuhanddienstleistungen am Markt
Kein Branchenausschluss (FR1.2)	erfüllt	siehe Abschnitt 2
Gesetzeskonformität (FR1.3, FR1.4)	vorausgesetzt	intern zu bestätigen
Transparenzbereitschaft (FR1.5)	zu entscheiden	siehe Abschnitt 5

Figure 16: The generated eligibility report. A results table for FR1, asset management classified as investment advising, and the self-flagged version caveat on V2.2. Anyone who wants to check it can find the source for every statement.

10 · LEVEL 3

When a bespoke application pays off.

Level 2 stops working where many people repeat the same task, where a state has to be maintained over months, or where outside parties need to enter data. Then you need an interface, and in 2026 that's buildable without a development team. The price is responsibility, and it's rarely factored in.

Where building your own tool makes sense in compliance

APPLICATION CLASS	WHY IT MAKES SENSE	ASSESSMENT
Policy diff: comparing policies against a new version	The comparison is deterministic, and a human can check the result in seconds. Addresses the 44 per cent who name policy management, and the free-text finding 'updating manuals and policies'.	Building it yourself is defensible
Training-record tracker	50 per cent name tracking training records. At the same time, the AI-literacy duty under Article 4 of the EU AI Act, in force since 2 February 2025, is itself an evidencing duty. The use case pays for itself twice over.	Building it yourself is defensible
Certification navigator	A public, versioned, tabular requirements catalogue. Errors stay checkable. See Case study A, solved there at Level 2.	Defensible with final human review
ESG and sustainability one-pager	Large volumes of documents, fixed standard values, a calculable result. See Case study B.	Defensible with final human review
Business partner due diligence / third-party due diligence	50 per cent name it. But: personal data, assessments with legal consequences, close to Annex III of the EU AI Act. A data protection impact assessment is essential.	Only with a DPIA and legal review
Regulatory watchlist	The highest need (81 per cent research manually), and at the same time the highest hallucination risk, in precisely the domain where it's most costly.	Only as a flag, never as advice

Figure 17: Prioritised by the ratio of benefit to risk, not by how appealing each one looks. For none of these classes is there a publicly documented, named case study of a self-built tool. The need and the task scope are evidenced; the ranking is our own assessment.




Figure 18: A benefit-risk matrix for compliance applications, following the logic of the effort-benefit matrix in the MMIND digitalisation guide. Notably, the survey's most-requested feature (AI-detected anomalies, 81 per cent) lands in the bottom-right box, with no audit trail.

11 · CASE STUDY B · LEVEL 3

A sustainability report in 45 minutes instead of three days.

Built with Hanno Konrad Anstalt, engineers, planners and surveyors. A prototype that produces an evidenced sustainability report from unstructured project documents.

45_{min.}

Processing time for a project with 1,200 to 1,500 pages of project documents

0.20_{to 0.50 CHF}

Model cost per run

>800_{pages}

of standards documents from KBOB, ecobau, SIA and SNBS used as reference

The task

SMEs in construction face pressure from two directions. Sustainability-reporting requirements are rising, and large clients with their own reporting duties pass that pressure straight down the chain. A report for a single building project tied up several days of a qualified project manager's time.

What was built

A web application that reads in quotes, invoices and specifications, extracts materials and quantities, matches them to CO₂-equivalents from KBOB, ecobau, SIA and SNBS standards documents, calculates the result, and checks it against sustainability criteria.

The same task structure (unstructured documents, a fixed body of standards, a calculable result, a citation requirement) applies to 84 per cent of respondents.

The decisions that mattered

1

Full source references

Every statement carries a document name and a verbatim quote. Without this decision, the report would be unauditible and worthless for compliance.

2

Clarifying questions instead of assumptions

Where information is missing or contradictory, the tool asks rather than estimates. In the prototype, for example, over a reinforcement quantity that looked implausible in the quote.

3

Versioning with an audit trail

Every version is saved with a timestamp, and earlier versions stay retrievable. A direct answer to Article 26 of the EU AI Act.

4

Human review as a mandatory step

The first report is explicitly a draft. In practice, only the second or third version reaches a standard fit for official use.




Figure 19: The overall score for a real project, here 74 points and Silver standard. Below it, each dimension carries a reasoned explanation with values and a reference to the applicable standard, for example '3-pane glazing Uw 0.87 W/m²K under SIA 380/1'. Where evidence is missing, the tool flags it, in this example for RC concrete certificates in in-situ concrete work.

What we learned, including the uncomfortable part

What worked

- A two-stage analysis: a fast model for summaries, a strong model for the final evaluation. Cuts cost and runtime substantially.
- Structured output to a fixed schema rather than free text. Without this, data extraction isn't reliable.
- Transparency builds trust faster than accuracy does. A visibly flagged doubt is more convincing than a suspiciously clean number.
- Folder discipline on the user's side: project data and standards documents kept apart. The same point as in Case study A.

What didn't work, or cost more than expected

- The first version was never the usable one. Iteration isn't clean-up work, it's part of the method.
- Browser-based storage lowers the barrier to entry, but it's no basis for multi-user operation or retention duties.
- The accuracy targets (at least 90 per cent for CO₂ calculations, under 5 per cent critical errors) are goals set during the test phase, not independently verified results. We present them as such.
- From the moment third parties use the application, obligations arise that nobody had budgeted for beforehand. See Chapter 13.

An honest look at the 80 per cent figure

The prototype is promoted with time savings of up to 80 per cent. That figure comes from our own measurement on real projects and hasn't been independently validated. It also applies to the drafting step alone, not the whole process. Human review remains, and isn't fully reflected in the calculation. We still quote the figure because it gets the order of magnitude right. But it's an experience, not proof. Anyone deciding within their own organisation should run their own comparison on two or three real cases before investing.

What this means for a compliance function

The prototype is the business case

This didn't start with a return-on-investment calculation. It started with a prototype tested on real project data. Only that showed whether the extraction would actually hold up. That order is cheaper than any feasibility study.

The citation requirement is the quality control

A model forced to back every figure with a source and a standards reference has less room to invent freely. What remains stands out on reading. Without that requirement, it doesn't.

Level 3 needs an owner

Once third parties use the application, it needs a named accountable person, a security review and a retention rule. Anyone who can't resource that is better off staying at Level 2.

The real payoff of Case study B wasn't the report. It was recognising that the task structure in compliance repeats itself: documents, a body of standards, a citation requirement, human sign-off. Build it once, and you can build it again.

12 · PROVIDERS

What matters is where your documents sit.

The market sorted itself along two axes in 2026. The first, how well a model produces code and text, is rarely the bottleneck for compliance tasks. The second is what decides: where your documents sit, physically and legally.

TOOL	CLASS	DATA RESIDENCY, EU / CH	RELEVANT FOR CH AND FL
Claude Skills, Cowork, Projects, Tasks	Folder work, workflows held in a file, scheduling	Cloud; check enterprise options	The most mature implementation of the workflow-in-a-file and folder-work model. No self-hosting.
OpenAI Codex	CLI, cloud agent	'Europe (EEA + Switzerland)', region selectable only at project creation, not changeable afterwards	One of only two providers in the field that name Switzerland explicitly. CLI under Apache 2.0. No on-premises option.
GitHub Copilot	Agent inside the development environment	EU Data Boundary incl. EFTA, Switzerland and Liechtenstein since 1 May 2026, a 10 per cent surcharge	Formally the best coverage. Whether the actual coding agent runs within that boundary is documented inconsistently. Get it confirmed in writing before deployment.
Mistral	Chat, CLI, open weights	EU headquarters, on-premises option	The only provider offering an EU base, on-premises deployment, an open CLI and open weights. Genuinely covers 'exclusively in-house'. On-premises pricing isn't public, and performance trails the leading field.
Kimi / Moonshot	API, open weights	Based in Beijing, no commitments	Only interesting via self-hosted open weights. Using the API means exporting data to China with no adequacy decision in place.
Open models, self-hosted	own infrastructure	full control	The only option that rules out the CLOUD Act. The performance gap is real, and operating overhead sits entirely in-house.

Figure 20: A selection, as at 21 August 2026, from public sources and vendor documentation. Prices and feature sets change monthly; treat this table as a checklist, not a shopping list. Full sources in Appendix B.

A decision rule that actually works

Rather than comparing providers, classify the data first. After that, the choice is clear in most cases.

DATA CLASS A Public Legislation, standards documents, association notices, published reports. Rule: any capable cloud tool is fine. This is where you start, and it covers most of Task 1.	DATA CLASS B Internal, no personal data Your own policies, process descriptions, anonymised metrics. Rule: a cloud service with EU or CH data residency, contractually guaranteed, no training on your data, retention periods agreed.	DATA CLASS C Personal data or data under professional secrecy Whistleblower reports, investigation files, client data, professional secrecy. Rule: self-hosted open weights, or a host with no US parent company. Data residency alone isn't enough here; see below.
---	---	---

Figure 21: Data classes as a selection filter. The practical consequence is a welcome one: Task 1 and much of Task 2 can be implemented in Classes A and B, meaning straight away and with no infrastructure project.

The finding that surprises most people

Under the guidance note issued by the Zurich cantonal data protection commissioner in February 2026, US authorities can access personal data under the CLOUD Act even where it sits in data centres in Switzerland or the EU/EFTA area, without the consent of the state where the data is held and bypassing mutual legal assistance. For professional, tax, welfare and victim-support secrecy, ancillary data must be protected too. For Data class C, that means the question 'is the data in Switzerland?' is the wrong question. The right one is 'whose law reaches the operator?'

Settle the data class first and you skip the provider debate. Start with the provider debate and you'll have it twice.

13 · LIMITS

Risks that limit deployment.

A compliance paper with no counterargument would be worthless. Nothing in this chapter contradicts Part III. Everything in it shapes how you go about it.

Hallucination in legal questions is not solved

The first pre-registered study of specialised legal AI tools (Stanford RegLab) found hallucination rates of 17 to 33 per cent, in tools with access to curated legal databases. A self-built tool running on a general-purpose model is unlikely to do better.

For calibration: as at 9 June 2026, 1,598 court cases worldwide have documented AI-fabricated citations, and the number is rising. Sanctions range from a few thousand to over US\$100,000, with suspensions in individual cases.

Consequence: use regulatory monitoring exclusively as a flagging function, with a mandatory source check. Never as legal advice, never without a citation.

Self-built applications are often wide open

A scan of 5,600 production applications built with AI found not a single one with CSRF protection, and none with standard security headers. A second study found around 5,000 applications with misconfigured access settings, roughly 40 per cent of them practically unauthenticated. Exposed data included hospital floor plans and sales records.

In addition, AI-generated code contained a security-relevant flaw in 44 per cent of generation tasks; the security success rate stood at 56 per cent, practically unchanged from the year before.

Consequence: never put a Level 3 tool into production without third-party security review. Treat prototypes as publicly reachable from day one.

Build it yourself, and you become a provider

The mechanism people miss: under Article 25 of the EU AI Act, you become the **provider** of a high-risk system if you put your name on it, make a substantial modification, or change its intended purpose so that it becomes high-risk. Article 26 requires operators to provide human oversight by competent staff, retain automatically generated logs for at least six months, and inform employees before deployment in the workplace. The AI-literacy duty under Article 4 has applied since 2 February 2025, even for purely internal use.

The timetable keeps moving

Under the political agreement on the Digital Omnibus of 6 May 2026, high-risk obligations were postponed: Annex III to 2 December 2027, Annex I to 2 August 2028. The transparency duties under Article 50 stay at 2 August 2026. **Importantly:** these dates only become binding once published in the Official Journal. Older analyses still cite 2 August 2026 for high-risk obligations. Planning around that is planning on outdated ground. For Liechtenstein, the regulation has also not yet been incorporated into the EEA Agreement.

14 · FROM NEED TO IMPLEMENTATION

Optimise first, then digitalise.

This approach isn't an invention for this paper. It's the digitalisation method from the MMIND guide, developed jointly with the Swiss Research Institute of Small and Medium-Sized Enterprises at the University of St Gallen, applied here to a core compliance process. The guide's founding principle holds unchanged: **optimise first, then digitalise.**

1

Define use cases and quick wins

Collect every candidate from day-to-day work and place it on the effort-benefit matrix: quick wins straight away, large projects planned, time-wasters left alone. For a compliance function, a third axis joins the two the guide normally uses: the data class from Chapter 12. A quick win in Data class C isn't a quick win.

Outcome: a prioritised list of three to five candidates, with one chosen as the first process. Duration: half a day.

2

Analyse the core process

Map the current state: who does what, with what, using which data, and where the medium breaks. Then define the target state against the guide's three questions: **Eliminate:** which step can be dropped? **Simplify:** which steps can be merged? **Standardise:** what does the single, consistent workflow look like? Only this third question actually makes the process automatable. A process that isn't standardised can't be written down as a workflow.

Outcome: current and target state on one page, with decision points and stopping points marked. Duration: one day.

3

Set up a folder with the relevant data

The step that gets underestimated, and in both case studies the one that decided the quality of the result. In practice: a source register with the closed list of permitted sources. Separate subfolders for your own documents and external standards. Access rights following the granular model that 66 per cent require. A data class set and documented for every folder.

Outcome: a working folder structure with a source register and a documented data class. Duration: half a day, plus the organisation supplying the material.

4

Set up the workflow that runs the process

The target process from Step 2 becomes a workflow held in a file, with the three non-negotiable properties from Chapter 8: a closed set of sources, a citation requirement, and a stopping point for human decision-making. Then run it twice on real cases, sharpen it, and set it up as a scheduled task.

Outcome: a live, repeatable process with a checkable result, plus a time comparison against the old workflow. Duration: one day, plus two test runs.

What you end up with

Four steps, four checkable artefacts. Not a concept paper, but a live process, and the ability to build the next one yourselves.




Figure 22: The four artefacts from the workshop. Artefacts 2 and 3 keep their value regardless of which tool gets used later. That's by design.

Who this works for

The method assumes no IT knowledge. Anyone who can draft a policy and describe a process can help shape Step 4. That's the break from earlier digitalisation projects.

A sensible line-up: the person accountable for compliance, someone from the affected business area, and someone with access to the data store. IT is kept informed, not commissioned, as long as you stay within Data classes A and B.

What we don't promise

No full automation. No replacement for legal advice. No tool that goes live without review. And no figure we haven't measured ourselves.

What we do promise: after four steps, a process that used to run by hand is running, and the organisation can set up the next one on its own. That's the difference between a tool and a capability.

Don't start with the most complex process, start with a quick win. Early, visible results matter more than a perfect first attempt. That, too, is a lesson from the digitalisation guide, and it holds up here as well.

15 · GOVERNANCE

The governance minimum before the first workflow.

53 per cent name a lack of in-house AI governance expertise as their biggest hurdle. This list isn't a governance framework. It's the minimum we won't start a workshop without.




Figure 23: The governance minimum. Points 3, 4, 5 and 7 are anchored in regulation; points 1, 2 and 6 are practical prerequisites.

APPENDIX A

Topline by question.

All figures n=32 unless stated otherwise. Multiple-choice questions are given as number of mentions, so totals regularly exceed 32. Conditional questions with a reduced base are marked as such.

QUESTION	DISTRIBUTION
Sector	Financial services 9 · Industry/manufacturing 9 · Healthcare 4 · Construction/real estate 3 · Public administration 3 · Retail 2 · Energy 1 · Other 1
Organisation size	more than 1,000 employees 16 · 300–1,000 7 · 1–49 5 · 250–299 2 · 50–249 2
Compliance culture	built into standard processes 14 · mostly reactive 6 · lived proactively 5 · box-ticking only 4 · a lived competitive advantage 3
Weekly time spent	5–19 h 12 · 20–49 h 12 · under 5 h 4 · 50–100 h 3 · over 100 h 1
Top compliance topics	Data protection 20 · corruption and bribery 17 · internal policies 17 · conflicts of interest 14 · internal controls 12 · supply-chain due diligence 11 · money laundering 10 · fraud 9 · competition law 8 · export controls and sanctions 8 · workplace safety 7 · ESG 7 · AI regulation 5
State of AI regulation	initial enquiries 12 · governance in place 9 · monitoring, no action yet 9 · not yet on the agenda 2
Hurdles in AI compliance	lack of in-house expertise 17 · classifying own systems 15 · unclear requirements 15 · documentation and transparency duties 13 · interplay with data protection 9 · other 4 · not yet relevant 2
Automatable activities top-3 selection	Reporting 19 · training administration 16 · internal investigations 15 · monitoring and oversight 14 · manual data entry 13 · regulatory monitoring 9 · business partner due diligence 8
Automatable processes	annual risk assessment 23 · regulatory monitoring 18 · compliance reporting 16 · training material 16 · due diligence 16 · training records 16 · gifts register 15 · policy management 14 · whistleblower reports 10 · KYC checks 10 · PEP screening 9 · internal correspondence 6
Regulatory tracking today	manual research 26 · trade associations 20 · in-house legal team 17 · external advisers 16 · specialised tools 2
Effectiveness measurement	no 18 · yes 14
Reasons against measuring base: the 18	no data foundation 15 · difficulty quantifying soft factors 13 · no suitable tool 11 · resources and time 8 · never requested 7 · other 1
Reporting frequency base: the 14	annually 6 · quarterly 6 · half-yearly 2
Expected core functions	Reporting 29 · ongoing monitoring 24 · risk visualisation 23 · warnings and alerts 21 · workflow automation 20 · document and policy management 18 · audit support 12
Warning triggers	AI-detected anomaly 26 · deadlines falling due 23 · specific matters 22 · thresholds exceeded 19
Preferred visualisation	traffic-light system 27 · heatmap 20 · trend lines 15 · text summary 10 · geo-map 8
Available data	unstructured documents 27 · emails 22 · spreadsheets 22 · structured databases 18 · log files 11 · don't know 2
Systems to connect	ERP 20 · CRM 19 · Microsoft 365 19 · HR 16 · finance/accounting 12 · GRC module 10 · legal tech 8 · core banking system 6
Type of integration	fully automatic via API 15 · don't know 5 · planned batch import 5 · no preference 4 · manual upload 3
Where processing happens	exclusively in-house 16 · hybrid 8 · external cloud acceptable 6 · no preference 2
Hosting	private cloud 11 · on-premises 7 · don't know 5 · SaaS 5 · no preference 4
Operating model	run in-house 13 · run by the tool provider 9 · no preference 5 · don't know 3 · third-party provider 2
Architecture	integrated platform with modules 16 · best-of-breed 7 · no preference 5 · don't know 4
Access	standalone web dashboard 14 · within Teams/SharePoint 6 · don't know 4 · within an existing GRC module 4 · as a Power BI report 4
Access model	granular 21 · multi-tenant 7 · simple 4
Willing to be contacted	yes 10 · no 22 · email supplied 11

APPENDIX B

Sources.

All web sources accessed on 21 August 2026 unless stated otherwise. [S] marks secondary sources with no traceable methodology. Read these figures with that caveat in mind.

Primary data

- 1 MMIND.ai × Bratschi, AI Compliance Survey 2026, n=32, fielded 18 March–31 May 2026.
- 2 MMIND.ai × Hanno Konrad Anstalt, prototype 'sustainability one-pager for construction projects', project presentation, November 2025.
- 3 MMIND.ai, B Corp eligibility report for a fiduciary firm in Vaduz, as at 20 August 2026, based on B Lab Standards V2.2 (20 February 2026).
- 4 MMIND.ai and the Swiss Research Institute of Small and Medium-Sized Enterprises, University of St Gallen, 'Digitalisation Guide for SMEs', 2025.

EU regulation

- 1 Regulation (EU) 2024/1689 (EU AI Act), in particular Art. 4, 25, 26, 50; EUR-Lex and the European Commission's AI Act Service Desk.
- 2 Gibson Dunn, 'EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines', on the political agreement of 6 May 2026. The postponement only becomes binding once published in the Official Journal.
- 3 Latham & Watkins and Mayer Brown on Art. 4 (AI literacy), in force since 2 February 2025.
- 4 Regulation (EU) 2024/2847 (Cyber Resilience Act); Directive (EU) 2024/1760 (CSDDD), including the 2025 Omnibus simplification package.
- 5 Lenz & Staehelin, on the extraterritorial effect of the AI Act for Swiss companies.

Swiss and Liechtenstein regulation

- 1 Federal Council, 'Position paper on the regulation of artificial intelligence', press release of 12 February 2025.
- 2 FDPIC, 'AI and data protection', as at 24 September 2025.
- 3 FINMA, Supervisory Communication 08/2024 of 18 December 2024.
- 4 Zurich cantonal data protection commissioner, guidance note 'Cloud products from US companies', February 2026 (CLOUD Act).
- 5 On the EEA status of the AI Act for Liechtenstein, and on the Liechtenstein government's AI strategy of 14 April 2026 [S], to be replaced with a government or EFTA primary source once published.

Security and error rates

- 1 Veracode, '2026 GenAI Code Security Report', 28 July 2026.
- 2 Cloud Security Alliance Labs, 'AI-Generated Code Vulnerability Surge 2026', 4 April 2026, updated 20 May 2026.
- 3 Cloud Security Alliance Labs, 'Vibe Coding AI Governance Gap', 2 June 2026, includes the IBM figures from 'Cost of a Data Breach 2025'.
- 4 IANS Research, 'Easy to Build, Easy to Expose', 15 May 2026.
- 5 Help Net Security, on prompt injection as the most common cause of agentic security incidents (OWASP), 11 June 2026.
- 6 The Hacker News, on the documented repository upload by a CLI version, 14 July 2026.
- 7 Gravitee, 'State of AI Agent Security 2026', n=919, reporting on a secondary source [S].

Hallucination and productivity

- 1 Magesh et al., Stanford RegLab, 'Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools', 2024. The oldest source in this paper; read accordingly.
- 2 Database of fabricated AI citations in court proceedings, as at 9 June 2026 [S], primary source Charlotin, HEC Paris.
- 3 METR, randomised controlled trial on developer productivity with AI assistance, 2025 (arXiv).

Tools and data residency

- 1 OpenAI, 'Your data', data residency 'Europe (EEA + Switzerland)', vendor documentation.
- 2 GitHub Changelog, 'Copilot data residency in US, EU and FedRAMP', 13 April 2026, updated 24 April 2026, EU Data Boundary incl. EFTA since 1 May 2026.
- 3 Mistral AI, product and licensing documentation for Le Chat, Vibe CLI and open weights.
- 4 Cursor, 'Models and pricing', vendor documentation.
- 5 xAI/SpaceXAI, 'Security FAQ', vendor documentation.
- 6 Moonshot AI, model and licensing details for Kimi K2.7 and K3 (first-party benchmarks, not comparable to SWE-bench Verified).
- 7 Anthropic, product documentation for Skills, Cowork, Projects and Tasks.

Standards and frameworks

- 1 ISO 37301:2021, compliance management systems.
- 2 NIST AI Risk Management Framework (AI RMF 1.0).
- 3 U.S. DOJ, 'Evaluation of Corporate Compliance Programs', September 2024 edition.
- 4 B Lab, 'Get started with the new standards', as at 21 August 2026.

APPENDIX C

How the free-text answers were coded.

Four questions were open-ended: pain points (Question 5), relevant areas of law (7), frameworks used (8), and desired features (24). Coding is interpretation. We set out the scheme here so the classification stays checkable.

CATEGORY	RULE	EXAMPLE FROM THE RESPONSES
Data, tools and reporting	Mentions pointing to a missing data foundation, unsuitable systems, or effort spent producing reports	'Systems that are hard to adapt technically, which adds manual work on top'
Information overload and regulation	Mentions of the volume, pace or multilingual nature of regulatory requirements	'Flooded with information and rules'; 'keeping an overview of current requirements, including in other jurisdictions'
Keeping policies current	Mentions of maintaining, updating and embedding a firm's own rulebook	'Updating manuals/policies in line with current developments in compliance'
Culture and integration	Mentions of how embedded the function is in the organisation, including negative self-descriptions	'Not integrated enough, lots of disconnected actions but no overall picture'; 'still too often seen as a necessary evil'
Effort on individual cases	Mentions of the effort spent on investigations, enquiries and false positives	'Too many false positive hits on our compliance filters, which then have to be cleared manually'
Training and evidencing	Mentions of training administration, evidencing and tailoring content to audiences	'Training and training records for mandatory compliance topics'

Notes on fair dealing

- **Multiple coding is allowed.** An answer touching several themes was assigned to several categories. Category totals therefore exceed the number of responses.
- **Blank and placeholder answers** such as 'none', 'n/a' or '–' were not coded; they were treated as missing and excluded from the relevant base.
- **Two answers on effectiveness measurement** cited reasons against measuring despite having answered 'yes' to the question. We applied those reasons only to the base of the 18 'no' answers; the raw total across all 32 would be one higher for some reasons.

Contact



Hartmut Hübner

Co-Founder

hartmut@mmind.ai

MMIND GmbH

Duxgass 55

9494 Schaan

Liechtenstein



Christian Wind

Partner

christian.wind@bratschi.ch

Bratschi AG

Bahnhofstrasse 70

P.O. Box

CH-8021 Zurich