

DATA AND CONTROL

Sovereign by Design

Keeping client data, AI models and control where they belong in regulated operations

Made by the team at Jaid

24 October 2026

Sovereign by Design

Executive summary

For years, data sovereignty in financial services meant one question: in which country is the data stored? AI has made that question far too small. **When AI reads client correspondence, the data also travels to wherever the model runs, may be kept by whoever runs it, and can end up shaping a model the firm does not own.** Most firms cannot draw that journey for a single client email.

That matters more in 2026 than it did two years ago. The legal basis for EU-US data transfers is under fresh pressure. DORA now requires EU financial entities to register every ICT third party, including AI providers and their subcontractors. And the clients of fund administrators, transfer agents and corporate service providers increasingly write their own data location and AI clauses into contracts.

This paper argues that sovereignty for AI in operations has four layers, not one:

1. **Data:** where client correspondence and case records are stored.
2. **Processing:** where the AI actually runs when it reads a request.
3. **Models:** whose model it is, and whether client data is used to train it.
4. **Control:** who can access, change, move or switch off any of the above.

A firm is only as sovereign as its weakest layer. A system that stores data in London but sends every email to a model in another jurisdiction, under someone else's retention policy, is not sovereign in any sense a client or regulator will accept.

We set out design principles for being **sovereign by design**, show how Jaid applies them, and close with questions any firm can ask of any AI vendor about its data.

1. Why sovereignty became an operations question

Sovereignty used to belong to the infrastructure team. Pick a data centre region, sign the right contract, tick the box. Operations teams rarely thought about it, because their tools did not move data anywhere new.

AI changes that, and operations centres feel it first. Their raw material is client correspondence: investor details, account numbers, transaction instructions, KYC documents, side letters. That is some of the most sensitive data a firm holds, and it is exactly what AI is now being asked to read.

Three things make this an operations question rather than an infrastructure one:

- **The data moves on every request.** A system that stores email in one place but sends each message to an external model for triage or drafting has created a cross-border transfer per email, not per contract.
- **The decision is made by whoever picks the tool.** When an ops worker pastes a client email into a personal chatbot, or a team wires a mailbox to a model API, a sovereignty decision has been made without anyone in risk or IT knowing.
- **Clients ask operations, not infrastructure.** Fund managers and their investors now ask their administrators directly: which AI touches our data, where does it run, and is our data used to train it? The answer has to come from the people running the service.

In short, every AI deployment in an operations centre is also a sovereignty decision. Firms that make those decisions deliberately can say yes to AI with confidence. Firms that do not will eventually be asked a question they cannot answer.

2. The four layers of sovereignty

To say a system is sovereign, a firm needs a clear answer at each of four layers. Each layer has its own failure mode, and a strong answer at one does not make up for a weak answer at another.

Layer	The question	What good looks like	Common failure
Data	Where are correspondence, attachments and case records stored, including backups and logs?	A region the client chose, stated in the contract, with backups and logs in the same place	Primary data in-region, but logs, backups or search indexes elsewhere
Processing	Where does the AI run when it reads a request, and where do prompts and outputs go?	Inference in the same jurisdiction as the data, or a documented, lawful transfer the client has approved	Data stored locally but sent to an external model API in another country on every request
Models	Whose model is it, and does client data train it?	No training on client data without explicit consent; client-specific learning kept separate from shared models	Prompts and documents retained by a model provider and used to improve a general model
Control	Who can access, change, move or switch off the system and its data?	Named access with full audit; the client can export, delete and exit; keys and configuration under clear ownership	Vendor support staff or sub-processors with broad, unlogged access; no tested exit route

The control layer is the one most often missed, and the one regulators care most about. DORA, for example, expects financial entities to understand their ICT providers' subcontracting chains and to have tested exit strategies. Knowing where data sits is not enough if you cannot show who can reach it, or how you would leave.

3. Where AI quietly breaks sovereignty


Few firms set out to send client data somewhere it should not go. It happens through ordinary decisions that each look reasonable on their own. Follow one client email through a typical AI-assisted operations centre and the leaks become visible.

1. **The personal assistant.** An ops worker pastes the email into a consumer chatbot to draft a reply. The client's details are now held by a provider the firm has no contract with, in a jurisdiction it has not assessed, under a retention policy it has never read.

2. **The model API call.** A sanctioned tool sends the email to an external model for triage. The firm's data stays in-region; the email does not. Unless the vendor can say where inference runs, the firm cannot either.
3. **The logs and telemetry.** Prompts, outputs and error traces are kept for debugging and monitoring, often by a different provider in a different region from the primary data. Logs are where sensitive data most often ends up forgotten.
4. **The training loop.** Some services use customer inputs to improve their general models unless the customer opts out. Once client data has shaped a shared model, it cannot be taken back out.
5. **The sub-processor chain.** The AI vendor uses a model provider, which uses a cloud provider, which uses its own support staff. Each link can access or move data, and each needs to be on the firm's register.
6. **The support session.** A vendor engineer investigating a problem is given broad access to production data from another country. It is rarely logged in a way the client could review.

FIGURE 1 / One email, six possible exits

POTENTIAL DATA EXITS



One email can cross the boundary through ordinary tools and defaults.

Each route needs an explicit decision about location, access, retention and control.

Every exit starts from an ordinary component inside the boundary; none of them requires a bad decision, only an unexamined default.

None of these is exotic. Each is a default somewhere. A sovereign design is one that makes the safe choice the default at every one of these points, and records that it did.

4. The rules that apply

There is no single sovereignty law. Instead, several regimes overlap, and an AI system in an operations centre usually touches all of them at once.

International data transfers (GDPR and UK GDPR). Personal data leaving the EU or UK needs a lawful transfer basis. For the US, most firms rely on the EU-US Data Privacy Framework, adopted in July 2023. It survived a first challenge at the EU General Court in September 2025, but that ruling is under appeal to the Court of Justice. In July 2026, following a US Supreme Court judgment on the independence of agencies like the FTC, the European Data Protection Board asked the Commission to examine whether the framework's foundations still hold, and privacy campaigners called for its withdrawal. Its two predecessors, Safe Harbour and Privacy Shield, were both struck down. A firm whose AI depends on routine EU-to-US transfers should have a plan for the day the framework falls.

Operational resilience (DORA). Since January 2025, EU financial entities must keep a register of every ICT third-party arrangement, including the subcontractors behind each provider, and submit it to their regulator every year. AI services, including calls to external model APIs, are ICT services in scope. DORA also expects due diligence before contracting and exit strategies for important providers.

UK outsourcing and third-party rules. UK firms face equivalent expectations through the FCA and PRA rules on outsourcing, operational resilience and the critical third parties regime.

The EU AI Act. For systems in scope, it adds obligations on data governance, logging, human oversight and cybersecurity.

Client contracts. Often the strictest regime of all. Fund managers, pension schemes and banks increasingly specify data location, sub-processor approval, AI use and training restrictions in their service agreements, and expect their administrators to evidence compliance.

The practical consequence is simple. For every AI component touching client data, a firm needs to be able to say where it runs, who provides it, who those providers rely on, whether client data trains it, and how it would replace it.

This section summarises regulatory direction for a general audience and is not legal advice.

5. Case studies: what others learned

Meta, 2023: when the transfer basis fails, the cost is operational as well as financial. In May 2023 the Irish Data Protection Commission fined Meta €1.2 billion, the largest GDPR fine to date, for continuing to transfer EU users' personal data to the US after the Court of Justice struck down Privacy Shield. The fine made the headlines, but the orders mattered more: Meta was told to suspend future transfers and bring its processing, including storage of EU data in the US, into compliance within months. The lesson for any firm relying on cross-border AI processing is that **a transfer mechanism can disappear, and the regulator may not give much time to adapt.** Systems designed to keep data and processing in-region do not face that cliff edge.

Samsung, 2023: sovereignty lost one paste at a time. In April 2023, engineers at Samsung's semiconductor business pasted confidential source code and internal meeting content into a public chatbot to get help with their work. Once submitted, the company had no way to retrieve or delete it.

Samsung responded by banning generative AI tools on company devices while it built internal alternatives. The lesson is twofold. **Sovereignty can be lost by well-meaning staff in seconds**, through exactly the personal-AI behaviour described in our companion paper *Off the Record*. And bans are a stopgap: the lasting fix was to give staff a governed tool that did the job without the data leaving the firm's control.

Together they mark the two ends of the risk: a systemic failure of the legal basis for transfers, and an everyday failure of individual behaviour. A sovereign design has to answer both.

6. Design principles for sovereign AI in operations

Sovereignty cannot be added after deployment. It has to be a property of the architecture, so that the safe choice is the default at every point where data could leak. We work to seven principles.

1. **The client chooses the region, and everything follows it.** Primary data, backups, logs, search indexes and AI processing all sit in the jurisdiction the client selected. No component quietly lives somewhere else.
2. **Processing goes to the data, not the other way round.** Where a model is needed, it runs in the same jurisdiction as the data it reads. Any exception is explicit, documented and approved by the client.
3. **Send the model only what the task needs.** A governed task retrieves only the fields and documents it is permitted to use. Less data in every prompt means less data to protect, everywhere downstream.
4. **Client data does not train shared models.** Learning from a client's examples improves that client's models only. Nothing a client sends is used to improve a general model without explicit, written consent.
5. **Models are replaceable.** No single model provider is a hard dependency. If a provider changes its terms, its location or its behaviour, the firm can switch, and the evaluation suite shows whether the replacement performs as well. This is what makes a DORA exit strategy real rather than theoretical.
6. **Every access is named and recorded.** Every human and system access to client data, including vendor support, is attributable and visible to the client. Production access for support is time-limited and approved.
7. **Exit is tested, not promised.** A client can export its data, cases and configuration in a usable form and have the rest deleted, and that process has been run, not just written down.

The last principle is the honest test of the others. A vendor that can show you a working exit has usually done the rest properly.

7. How Jaid approaches sovereignty

Jaid serves operations centres on both sides of the Atlantic, from offices in London, Cambridge and New York. Our clients include fund administrators, prime brokers and corporate service providers whose own clients set strict terms on where their data goes. Sovereignty is therefore a design requirement for us, not a feature.

Principle	How Jaid applies it
Client-chosen location	Colocation in the UK, the EU or the US, chosen by the client. Correspondence, case records, backups, logs and AI processing all stay in that location
Processing goes to the data	AI models run in the same location as the data they read
Minimum data per task	Every AI task is a governed task with a defined scope. It retrieves only the evidence its permissions allow, so each model call carries the least data needed
No training of shared models on client data	Jaid shapes intelligence around each client's own terminology, request types and rules. Client data is never used to train models shared with other clients
Replaceable models	We compare a range of open-weight models and closed-source models agreed with the client, workflow by workflow. Evaluation suites show whether a replacement performs as well, so switching is a measured decision, not a leap
Named, recorded access	Every access to client data is recorded, alongside every request, retrieval, AI action and human decision on the case
Tested exit	Clients can take their data, cases and configuration with them and have the rest deleted

The common thread with our other papers is the case record. Because every step of the work, human or AI, is recorded on the case, the same record that proves how a request was handled also shows which components touched the data and where. That is what lets a client answer a sovereignty question from evidence rather than from a contract clause.

8. Questions to ask any AI vendor about your data

Use these with any vendor whose AI will read client correspondence, including us. Ask for evidence, not assurances.

Data

- ☐ In which jurisdiction are our correspondence, attachments and case records stored, including backups, logs and search indexes?
- ☐ Can we choose that jurisdiction, and is it written into the contract?

Processing

- ☐ Where does the AI model run when it reads one of our emails, and does any of our data leave that jurisdiction at any point?
- ☐ How much of each email and attachment is sent to the model, and why?

Models

- ☐ Is any of our data, including prompts and outputs, used to train or improve a model that other clients use?
- ☐ If you learn from our examples, where is that learning kept, and can it be deleted?
- ☐ If your model provider changes its terms or location tomorrow, how quickly can you switch, and how would we know the replacement performs as well?

Control

- ☐ Who at your company, and at each of your sub-processors, can access our data, and is every access logged where we can see it?
- ☐ Show us an exit: what we would receive, in what format, and how deletion is confirmed.

A vendor that answers these clearly, with documents and a demonstration, has designed for sovereignty. A vendor that answers with a policy statement has not.

Conclusion

AI has turned data sovereignty from a hosting decision into an operations decision made thousands of times a day, every time a client email is read. Firms that treat it as a box ticked at contract signature will find their data in places they never chose: a personal chatbot, a model provider's logs, a shared model, a sub-processor's support desk.

The alternative is to be sovereign by design. Keep data and processing in the jurisdiction the client chose. Send the model only what the task needs. Never train shared models on client data. Keep models replaceable, access recorded and exit tested. And keep a case record that shows, for every request, exactly what touched the data and where.

That is how we build Jaid. We would welcome the chance to walk your operations, technology and risk teams through the journey of one of your client emails, layer by layer.

Sources

- [Computing: US Supreme Court ruling casts doubt on EU-US data privacy deal](#) (2026)
- [PPC Land: EDPB puts EU-US data pact in doubt](#) (August 2026)
- [WilmerHale: Court of Justice to review challenge to the EU-US Data Privacy Framework](#) (December 2025)
- [Sidley: Time to register your ICT third-party service providers under DORA](#) (April 2025)
- [Infosecurity Magazine: Meta fined €1.2bn for violating GDPR](#) (May 2023)
- [Fortune: Samsung bans employee use of ChatGPT after data leak](#) (May 2023)

Contact: success@jaid.io