CHAPTER 4

Who Is Liable — compliance as organizational architecture

> *The EU AI Act, GDPR and their American counterparts turn liability into a question of construction. An organization that ignores the question builds a risk and calls it an AI system.*

Regulation is the legal projection of the gap between machine generation and human verification. The EU AI Act, GDPR and their counterparts across the Atlantic leave cognitive asymmetry fully in place — they fix it in law. The gap between what a model generates and what a human manages to check gets translated into the language of rights and duties: who is the Provider — under the EU AI Act, the party that develops an AI system or places it on the market; who is the Deployer — the party that uses the system in its own operations; who is the AI operator — the person who runs the system day by day; who is the Compliance Officer — the person who answers for regulatory adherence. Compliance is the same asymmetry, discovered by lawyers.

The EU AI Act entered into force in August 2024. The transition periods for most high-risk categories have moved into their active phase or are entering it. Most machinery manufacturers in the DACH region have yet to feel the regulation's full weight. Once the transition periods lapse, the rules bind — with fines that, for the gravest violations, reach up to 35 million euros or 7 percent of global annual turnover, whichever is higher.

The question *who is liable* belongs to construction rather than to legal post-processing. Most projects ask it too late — at the moment the architecture window has already closed. Asked on time, it becomes a construction discipline with a small set of concrete artifacts: a classification, a form, a matrix, an insurance line — and the answer stays load-bearing when it matters.

How Europe and the United States approach it

The United States has no single federal counterpart to the EU AI Act. Regulation splits across three levels: the NIST AI Risk Management Framework — a voluntary methodological standard from the US standards institute; sector rules — the Federal Trade Commission for advertising claims, the FDA for medical AI, the Equal Employment Opportunity Commission for hiring; and state laws such as the Colorado AI Act of 2024 or California's SB 1047. The overall logic — risk categories, documentation, audit, disclosure duties — sits close to the European one. The differences lie in the speed of entry into force, the level of sanctions, and how heavily the duties concentrate on the Deployer.

The European regime is the stricter one, and that settles the practical sequencing. Whoever builds compliance under the EU AI Act covers, in the great majority of cases, the American requirements as well. The reverse route — starting from the laxer regime and upgrading on entry into the EU market — usually costs more, because it forces the rework of architectural decisions taken without the EU risk categories in view.

The EU AI Act in practice: high-risk classification, documentation duties, named persons

The regulation names a concrete list. For an industrial DACH context, three categories matter most.

Product safety. AI systems that serve as safety components of products already governed by existing EU regimes — the Machinery Regulation 2023/1230, the Medical Device Regulation, the aviation safety rules. If an AI system influences the safety of the end product in any way, it very likely falls under high risk.

Critical infrastructure. AI systems that manage water supply, energy or transport. For industrial manufacturers that form part of critical infrastructure, this brings additional requirements.

Employment and workforce management. Systems for recruiting, performance evaluation, task allocation or monitoring. Every AI system that touches decisions about employees deserves careful examination.

For high-risk systems the EU AI Act requires: a risk-management system before deployment, data-quality management for the training sets, technical documentation under Annex IV, automatic logging, transparency toward end users, human oversight — the point where the 30-percent layer starts to work — and accuracy, robustness and cybersecurity.

And, critically for the argument here, a clear definition of *Provider* and *Deployer*. The Provider develops the system or places it on the market. The Deployer puts it to use. Both carry duties. A contract may distribute liability between them — the regulator looks at both sides regardless.

Stuart Russell, in *Human Compatible*, states the general rule: liability for an autonomous system must have an institutional bearer before the system reaches the market.[1] Kate Crawford adds, in *Atlas of AI*, that this bearer is never a *purely* technical role — it is always embedded in a regulatory, contractual and organizational architecture.[2] Read in that frame, the EU AI Act writes into law a requirement that already existed: between the model and the market stands a named person with a mandate.

The liability form: five questions for every AI process

The liability form — one AI process, one page

1. Classification

EU AI Act: high risk (Art. 6 + Annex III) / limited / minimal · rationale · signature

2. Named persons

Provider: _____ Deployer: _____

AI operator: _____ Compliance Officer: _____

3. Human intervention points

Escalation thresholds · who authorizes an override · who logs the decision

4. Documentation

Annex IV · logs · risk-management file · incident register

5. Insurance coverage

Policy · limits · exclusions · who holds the claim after an incident

The structure of the liability form for one AI process — one workload, one product stream. Fill-in order: (1) classification under the EU AI Act — at the end of the pilot, before the decision to go productive; (2) named persons — at the planning stage, before development starts; (3) human intervention points — at the design stage, before training and integration; (4) documentation — grows step by step, finalized before launch; (5) insurance coverage — before launch. Every empty block works as a hidden risk: liability exists structurally only where every corner has a named bearer with a written mandate, and an empty corner is a structural deformation that surfaces at the moment of an incident. Adapted from the EU AI Act Implementation Toolkit of the European Commission's Joint Research Centre[3] and the RACI schemas of ISO/IEC 42001, the AI management system standard.[4]

Distributing liability by module: who signs what, when, with which insurance

In real projects liability is rarely monolithic. A typical AI system consists of several components, each with its own chain of liability.

Model. Who is the Provider? For a bought model — the vendor. For a model trained in-house — the company itself. For a third-party model that has been fine-tuned — further trained on the company's own data and context — liability is split, and the details hang on the license contract.

Infrastructure. Where does the model live? On-premises, the company answers for availability and security itself. In the cloud, responsibility is shared — and the Deployer duty under the EU AI Act shrinks by nothing.

Integration. Who wired the model into the production environment? An external integrator? The in-house IT team? What integration documentation exists?

Operations. Who runs the system day by day? That is the AI operator — named, by name, in the compliance documentation.

The practical tool is a module-by-module responsibility matrix, a RACI variant for AI: Responsible — who does the work; Accountable — who answers for the result; Consulted — who advises; Informed — who is kept in the loop. Each row documents, for one module, who is Responsible, who is Accountable, who signs, and which insurance covers it. This is no bureaucratic end in itself. The matrix answers the question *who signs* at the moment something has gone wrong. Something always goes wrong.

RACI matrix for an AI system: one row per module

Module	R	A	C	I	Insurance
Model	Data team	CTO	Compliance	CEO	Cyber + E&O
Training data	Data curator	DPO	Legal	CTO	GDPR liability
Infrastructure	IT ops	CTO	CISO	CFO	Cyber policy
Integration	Integrator	Head of IT	QA	CTO	Vendor SLA
Operations	AI operator	Process owner	Compliance	CEO	D&O + E&O

R — does the work · A — answers for the result (one per row) · C — consulted · I — informed

A filled-in RACI matrix for a typical high-risk AI system in an industrial DACH context. Each module — each row — has exactly one Accountable; the other roles follow actual competence. The insurance column names the policy type covering incidents in that module: Cyber + E&O (errors and omissions), GDPR liability, D&O (directors and officers), vendor SLA (service level agreement). An empty cell is a structural defect. ISO/IEC 42001 extends classic RACI with the duty to document the names of the role holders, not only the job titles.[4]

Sovereign-washing: why an “EU cloud” rarely releases anyone from liability

Sovereign-washing — a feigned sovereignty — is a marketing practice spreading through the DACH market. Cloud providers advertise a “GDPR-compliant EU cloud” or “sovereign AI” and imply that these labels lift the liability for the system's results off the customer.

They do nothing of the kind.

The Deployer duty under the EU AI Act is independent of geography. Whoever deploys a high-risk AI system carries the Deployer duties — in a data center in Munich, in Paris or in Dublin alike. An “EU cloud” means the data is stored inside the EU. That can ease GDPR compliance for data transfers. It changes nothing about who answers for the system's results.

A data processing agreement — the GDPR contract, in German practice the Auftragsverarbeitungsvertrag or AVV — makes the cloud provider a processor acting on the customer's instructions. When something goes wrong, the regulator turns first to the customer as the data controller. A recourse claim against the provider may exist — and pursuing it is the customer's problem, never the regulator's.

Availability is a separate risk. An “EU sovereign cloud” often means regional infrastructure controlled by a US parent company. In a scenario of geopolitical tension or widened sanctions, that “sovereign” cloud can become unreachable for legal reasons while every server keeps humming.

> *No certificate, no data processing agreement, no marketing label about “sovereignty” or “EU compatibility” replaces a compliance architecture of one's own, with clearly named responsible people.*

The anthropological point

Compliance is often perceived as a constraint — something bolted on top of the system to satisfy the regulator. That perspective hides its real function. Compliance is the structural layer that exposes the liability decisions already present in the system. A missing responsibility matrix leaves the liability fully intact; it merely leaves it unnamed — and at the moment of an incident the liability lands where nobody expects it. Usually on the weakest link of the organizational hierarchy.

Conclusion

The architectural approach to compliance means constructing the organizational layer of the system with the same seriousness as its technical layer. Named responsible people, written mandates, insurance coverage, and a liability form the AI operator uses as a daily working document — a living page rather than a file in the legal department's archive. Where all of this exists, compliance stops being a risk and becomes a support. Where it is missing, any technical architecture stands on soft ground.

Three Questions for Your Own Context

1. **Classification.** Which of your AI processes — individual workloads — fall under high risk under the EU AI Act? Does each one have a written classification with a rationale — and who signed it?
2. **Distribution by module.** Does a responsibility matrix exist that names, for every module — model, data, infrastructure, integration, operations — the Provider, the Deployer and the insurance coverage? If it does not: which modules are the most critical?
3. **The sovereign-washing test.** Which compliance claims of your cloud or AI vendors rest on marketing labels rather than on contractually fixed, verifiable duties? What happens to your liability if their "EU cloud" becomes unavailable?

Sources

1. Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control*, Viking, 2019.
2. Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*, Yale University Press, 2021.

3. European Commission, Joint Research Centre, *EU AI Act Implementation Toolkit*, 2024.

4. ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*, 2023.