Assigning a responsibility does not guarantee the means to exercise it

Forward Deployed Engineers and regulatory continuity: read the job postings, the terms and the law before presuming

Several major players invested heavily, in the spring and summer of 2026, in the embedded deployment of AI in the enterprise: AWS announced a one-billion-dollar investment in a *Forward Deployed Engineering* organisation, Microsoft launched Frontier Company, an operating business of some 6,000 experts backed by a 2.5-billion commitment, and OpenAI set up a deployment company funded with more than 4 billion by third-party investors. The three amounts do not measure the same thing; they indicate a direction, not a market norm. The model is well known: the vendor's engineers build in the client's environment, deliver within a few weeks, then leave. What remains is the aftermath. Who then holds the regulatory obligations attached to the system, and with what means?

The thesis is narrow. An FDE contract can clearly assign operational and compliance responsibilities to the client without describing, in the same document, the information, access and assistance useful for exercising them. *The allocation of obligations and the provision of means are therefore to be verified separately, system by system*, taking into account the actual legal role of each actor. The domain of validity is that of custom-built systems likely to fall under the high-risk regime of Regulation (EU) 2024/1689. Three materials are used, and none draws conclusions in place of the others: job postings, public contractual terms, the law.

1. What the job postings make visible

Twenty-two public FDE management job postings were coded, consulted on 13 November 2026: fourteen at hyperscalers and laboratories (Google Cloud, Microsoft, AWS, OpenAI, Anthropic, Scale AI, Cohere, Snorkel, Handshake), eight at platforms or AI application vendors. The inventory, the links and the codebook are given in Appendix A. According to this codebook, which counts as regulatory only a mention naming an external obligation or standard, two postings out of twenty-two explicitly mention a compliance standard; none cites the AI Act; none of the seven hyperscaler postings contains a mention of this type. All twenty-two, by contrast, speak of adoption, growth or speed.

The two occurrences are instructive. At Scale AI, compliance (SOC 2, GDPR, HIPAA) appears as a boundary to be respected within the client's security and compliance perimeter; at Parloa, it takes the form of expected experience with the GDPR and the ISO 27001 standard. In both cases, the subject is data protection and information security, not a regulation specific to AI systems. Elsewhere, the governance terms that appear in the postings relate to security, data sovereignty or responsible AI. According to the codebook adopted, mentions of explicit regulatory compliance therefore remain rare in the corpus; this describes the vocabulary of the postings, not the place of compliance in the actual organisation.

The result measures visibility, not organisation. A posting indicates what the employer chooses to put forward in order to recruit; it says nothing about internal objectives, delivered artefacts or the teams that, under other job titles, carry compliance. The corpus is small, non-random and mostly American; it was coded twice by the same coder, but neither by a second coder nor against a control group covering the compliance postings of the same companies, which would make it possible to measure the visibility of those functions. The chain that would lead from "few mentions" to "the client inherits a debt" has four links; the corpus documents only one. Not knowing does not authorise presuming; it obliges one to read the contracts.

2. Technical run, provider monitoring, deployer monitoring

The same corpus qualifies a received idea: expectations of production, evaluation or reliability are explicit in nineteen postings out of twenty-two. What the postings make visible on the operational side is of a different nature from what they leave barely visible on the regulatory side.

For high-risk systems, and from 2 December 2027 for those in Annex III or from 2 August 2028 for those in Annex I, the regulation imposes two forms of monitoring that do not overlap. Article 72 requires the provider to establish a post-market monitoring system, based on a plan that forms part of the technical documentation. Article 26 requires the deployer to use the system in accordance with the instructions for use, to monitor its operation and, where relevant, to inform the provider in accordance with Article 72; it also requires the deployer to keep the logs automatically generated by the system, to the extent that such logs are under its control, for at least six months, unless provided otherwise in applicable law, in particular on the protection of personal data. Supervising a system so that it works and being able to establish that it remains compliant are two distinct activities, carried by distinct actors. *Technical run is not regulatory run*.

3. Reading FDE terms: a bounded case

AWS publishes the terms of its FDE offering, in a *preview* version updated on 12 November 2026. The document is public, precise and dated. It has three boundaries. It is a single vendor. These are *preview* terms. They refer to a wider set: the AWS Customer Agreement or another applicable agreement, a sprint letter that triggers the engagement, a security plan for access to production, addenda on regulated data and on health data, usage policies. What follows describes the FDE terms consulted, not the contractual whole.

These terms first circumscribe the authorised uses. They exclude hazardous environments and critical systems; they are not designed for finished medical devices and do not authorise the solutions, on their own, for clinical decision-making; access to protected health data requires a signed addendum. The case therefore says nothing about how one contracts an FDE for a high-risk system in healthcare or critical infrastructure. It shows how a vendor manages its own risk: by delimiting what it authorises.

Within this perimeter, the allocation is explicit. The client decides on deployment, operates and maintains; it answers for activities and modifications in its production environment, including those carried out by AWS on its instructions; it manages access, logs and records; it ensures that the use complies with applicable laws, the AI Act being cited by name. The regulation absent from the postings thus appears in the contract, on the side of the client's obligations. Nothing more will be drawn from this.

The means provided are real. The client owns the content developed and delivered during the sprint (documents, software, models, prompts) and receives a perpetual licence to any pre-existing AWS intellectual property included in it; its engineers work alongside the FDEs during forty-five-day sprints. The terms consulted do not, on the other hand, describe any transfer of skills or any support after the engagement; they provide for the revocation of AWS's production access at the end of the work and exclude any legal or compliance advice.

Concluding that means are lacking would be a slippage: the absence of a promise in this document proves neither its absence elsewhere nor the absence of skills on the client's side. What the case establishes is narrower. The allocation and the ownership can be read in these terms; the sufficiency of the means for a given system cannot. Holding the code and the prompts does not say whether one holds the technical documentation, the evaluations, the logs and the assistance that the applicable obligations require. *Assigning a responsibility and providing the means to exercise it are two distinct verification points.*

4. Who is the provider? Facts, not a clause

The regulation defines the provider as the person that develops an AI system or has an AI system developed and places it on the market or puts it into service under its own name or trademark (Article 3, point 3). Article 25 then covers three situations in which a deployer or a third party becomes the provider of a high-risk system: putting its name or trademark on a high-risk system already placed on the market; making a substantial modification to such a system, provided that it remains high-risk; modifying the intended purpose of a system which, as a result, becomes high-risk. In these cases, the initial provider is no longer considered the provider of that system for the purposes of the regulation; Article 25(2) provides that it then cooperates closely with the new provider and makes available the information, the technical access and the assistance reasonably expected, unless it has clearly specified that its system is not to be changed into a high-risk system. The question of means therefore also arises at the moment when the role of provider changes hands.

A system custom-built by an FDE team can therefore make the client a provider; it does not necessarily do so. The decisive facts are identifiable: who develops or has the system developed, who puts it into service and under what name, for what intended purpose, with what modifications. The AWS terms provide one indication, since they prohibit marketing the solutions under the AWS brand; they do not settle the qualification and do not release any operator from its legal obligations. *A clause allocates commitments. It does not qualify a provider.*

5. The transitional regime: what is clarified, what is not

Since Regulation (EU) 2026/1744, in force since 27 July 2026, Article 111 specifies the treatment of high-risk systems already on the market. One part is clarified: these systems are concerned only in the event of a significant change in their design from the date of application; where a first unit was lawfully placed on the market before that date, other units of the same type and model may continue to be placed on it as long as the design remains unchanged. Two specific regimes are added: compliance by 2 August 2030 at the latest for systems intended to be used by public authorities, and a regime of its own for components of certain large-scale European information systems.

One part remains uncertain, and it is narrower than is sometimes said: not the existence of a framework, which the amending regulation has specified, but the practical application of the notion of a significant change in design to a new model version, to modified prompts or to a reworked architecture. On this precise point, no established doctrine has been identified to date [SOURCE TO BE DOCUMENTED]. A model update therefore does not automatically trigger the obligations. But a system delivered in iterations, dependent on components whose versions evolve, raises a practical question

that the calendar does not settle: who assesses the scope of each change, according to what criteria, with what information?

6. What the buyer should verify

Three checks, system by system. On allocation: the terms and the documents they incorporate, starting with the scope of authorised use. On means: what is provided, then what is sufficient in view of the role actually held (see below), including the transfer of skills, support after the sprint and the procedure for documenting changes. On role: the qualification in light of the facts, independently of the contract. The risk of a responsibility without sufficient means is to be assessed in each engagement; it is not a demonstrated consequence of the FDE model.

What one must possess depends on the role, and the obligations of one actor do not, on their own, create a general right of access to the documents of another. For a high-risk system, the technical documentation is drawn up and kept by the provider (Articles 11 and 18); it includes in particular the elements of risk management and testing, and a deployer does not, by virtue of its role alone, have a right to obtain it in full. What the provider must give the deployer is a set of instructions for use (Article 13). Logs are generated automatically by the system, whose design is the provider's responsibility (Article 12); each party keeps those under its control (Articles 19 and 26). Where the client is itself the provider, it must draw up the technical documentation and needs, for that purpose, elements that it does not produce alone.

This is where the cooperation obligations come in. In addition to Article 25(2) mentioned above, Article 25(4) requires the provider of a high-risk system and the third party that supplies an AI system, an AI model, tools, services, components or processes used or integrated in that system to specify, by written agreement, the information, capabilities, technical access and other assistance necessary to enable the provider to comply with its obligations. The exception covers third parties making publicly available, under a free and open-source licence, tools, services, processes or components other than general-purpose AI models. Support after the sprint and the transfer of skills are therefore not general, self-standing obligations owed in every FDE contract. But assistance after the sprint may fall within this written agreement where it is necessary to enable the provider of the system to fulfil its obligations. Beyond what is necessary in this sense, it is a matter of good procurement practice.

Limitations

The corpus measures visibility in twenty-two postings, thirteen of which show no date, with neither a second coder nor a control group. The contractual case is unique, in *preview*, and excludes several critical and clinical uses. The client-specific agreements,

sprint letters included, were not consulted. The notion of a significant change in design applied to AI models remains without established doctrine. The routes to sanction and redress, summarised in Appendix B, depend on conditions and timelines of their own that this text does not develop.

Conclusion

The FDE model is neither a threat to compliance nor a solution to its problem. It shifts the question towards a document that the job postings do not illuminate: the contract, read line by line, then confronted with the law and with the facts of the system. An allocation can be perfectly explicit without the means of exercising it being described in the same place. This is not a defect specific to one vendor; it is a property of any allocation.

The posting shows what is put forward. The contract organises what the parties owe each other. The law determines, case by case, who answers for what.

Appendix A. Corpus method

Selection. FDE management job postings (Manager, Senior Manager, Head, Director, Multidisciplinary Manager FDE) visible in 2026 on careers pages or aggregators; consulted on 13 November 2026. Four postings excluded (empty page, position filled, inaccessible page, uncodable text).

Inventory. The 22 postings are listed in corpus.csv, which accompanies the article: company, job title, segment, URL, displayed date, dates of consultation (13 November 2026) and recoding (14 November 2026), coding by category, coded excerpts. The excerpts, of ten words at most, are grouped by posting in fiches-corpus-2026-11-14.md. Dates: seven postings display an absolute date (from June to September 2026, one of which is an update date), two a relative date, thirteen none. Public archiving of the pages could not be carried out from the working environment [SOURCE TO BE DOCUMENTED].

Codebook. Unit: a sentence segment expressing an expectation of the role. Four categories fixed before reading: regulatory compliance (a named external obligation or standard, for example SOC 2, GDPR, HIPAA, AI Act, sector regulation); technical governance (security, data sovereignty and governance, responsible AI, guardrails); run and stability (production, evaluation, monitoring, reliability); commercial and speed (adoption, growth, expansion, revenue, time-to-value, speed). Exclusions: HR legal notices and mission statements. Coding examples: "within customer security and compliance boundaries (SOC 2, GDPR, HIPAA)" coded as regulatory compliance; "secure governance" coded as technical governance. The rule tends to underestimate the regulatory category.

Procedure. Automatic extraction of occurrences followed by manual review by a single coder, in two passes one day apart. Agreement between passes: 86 identical decisions out of 88 (22 postings, 4 categories). The two discrepancies are documented in corpus.csv: a compliance mention missed in the first pass (Parloa), a security clearance mention set aside in the second (AWS). The second pass is authoritative. Double coding by a second coder not carried out. Borderline case recorded: one posting (Anthropic) mentions intellectual property and contractual matters, not coded as regulatory compliance.

Results by posting. Explicit regulatory compliance: 1 of 14 hyperscalers and laboratories, 1 of 8 platforms. Technical governance: 7 and 1. Run and stability: 14 and 5. Commercial and speed: 14 and 8. The numbers of occurrences depend on the codebook and on the length of the postings; they are not used in the text.

Appendix B. Possible legal routes, subject to conditions

Market surveillance (Regulation (EU) 2024/1689, Article 79): for a system presenting a risk, corrective measures first; restriction, withdrawal or recall are among the possible measures. *Fines* (Article 99): ceilings, notably 15 million euros or 3% of worldwide annual turnover for most operator obligations, with assessment criteria and specific rules for SMEs. *Liability for defective products* (Directive (EU) 2024/2853): a regime provided for products, software and AI systems included, placed on the market or put into service from 9 December 2026, subject to national transposition; no-fault liability, but proof of damage, defect and causal link is required, subject to the presumptions; the manufacturer may be the developer or the manufacturer of a component. *Representative actions*: Directive (EU) 2020/1828, whose Annex I includes the AI regulation as a result of Article 110 of Regulation (EU) 2024/1689; harm to the collective interests of consumers. *Group action in France* (Law No. 2025-391 of 30 April 2025, Article 16): authorised claimants, persons in a similar situation, a common breach. The actors targeted and the conditions of liability vary by regime: the operator concerned (provider, deployer or other) under the AI Act, on the basis of its own obligations; the manufacturer under product liability, on proof of a defect, damage and a causal link, without any regulatory breach being required; the professional concerned in a collective action, depending on the legal basis invoked.

Main sources.

Regulation (EU) 2024/1689 and Regulation (EU) 2026/1744 (EUR-Lex); AWS, *FDE Preview Terms* (aws.amazon.com/legal/fde_terms, version of 12 November 2026); announcements by AWS (aboutamazon.com, 30 June 2026), Microsoft (blogs.microsoft.com, 2 July 2026) and OpenAI (May 2026, URL of the primary announcement [SOURCE TO BE DOCUMENTED]); Directive (EU) 2024/2853; Law No. 2025-391 of 30 April 2025.