

# The Contractualised Promotion Port

What the FDA Built Before the Architects Named It

Jérôme Vetillard  
Twingital Institute

April 24, 2026

*Why the Predetermined Change Control Plan (PCCP) is the mature form of an artefact  
that mainstream MLOps still confuses with the model registry*

---

## Contents

|  |    |
|--|----|
| 1. Introduction  | 1  |
| 2. Registry, Port, Promotion: What the Terms Confuse                 | 2  |
| 3. Diagnosis: The MLOps Misunderstanding                             | 3  |
| 4. Anatomy of the PCCP: The Contractualised Promotion Port           | 4  |
| 5. The Hexagonal Prerequisite and the Discipline of the Silent Trial | 6  |
| 6. Objection Addressed: Is the PCCP Transposable Outside Health?     | 7  |
| 7. ToxTwin as an Implementation Terrain                              | 8  |
| 8. Doctrinal Consequences  | 10 |
| 9. Conclusion  | 11 |

---

## 1. Introduction

On 14 January 2026, the FDA and the EMA jointly published ten *Guiding Principles of Good AI Practice in Drug Development*. On 2 February 2026, the US Quality Management System Regulation, aligned with ISO 13485, came into force for medical devices. Between 2025 and 2027,

the European regime applicable to general-purpose AI models under the AI Act progressively enters its operational phase.

These three milestones do not belong to the same legal regime. The FDA/EMA Guiding Principles frame the use of AI in the drug development cycle. The QMSR frames the quality systems of medical devices. The AI Act introduces a regime of its own for AI systems and general-purpose models. Conflating them would produce exactly the kind of regulatory homogenisation that serves to appear compliant without committing to actually being so. The convergence is functional, not legal.

This functional convergence is nonetheless real, and that is what concerns the architect. Three distinct regulators are shifting the same central question: it is no longer enough to know whether a model works, nor even whether it has been validated. One must be able to state *under what conditions* it may enter a regulated usage regime, *by what method* it may evolve, *with what traceability*, *under whose responsibility*, and *under what conditions* its admission may be withdrawn.

In a previous article, I posited hexagonal architecture as the structural frame of a pharmacovigilance twin. That article showed that the separation between business core and adapters is the condition under which a system remains auditable, testable, and substitutable on its peripheries. It did not settle the following question: *where, within this architecture, does the artefact stand that institutes the passage into regulated production?* That question is the object of the present text.

**My thesis.** The Predetermined Change Control Plan required by the FDA for AI-enabled medical devices is the mature form of what I propose to name the *promotion port*, that is, an ex ante instituted contract that distinguishes the regulated boundary from the versioned registry. So long as industrial AI confuses what medical device regulation has built with the model registry it has tooled, it will remain unable to durably cross the threshold of regulated production.

**Domain of validity.** The thesis holds for AI systems deployed in regulated production or candidate to regulation (health, finance, defence, automotive, pharmaceutical industry, agri-food, critical infrastructures). It does not apply as such to non-decisional internal uses (productivity assistants, non-exposed documentary tasks) where the model registry may suffice, the promotion instance being in that case degraded. But as soon as a system produces an inference that affects a decision, a priority, a classification, an investigation, an alert, an allocation, or a human trajectory, the registry no longer suffices. A port is required.

## 2. Registry, Port, Promotion: What the Terms Confuse

Three terms circulate in the MLOps literature: registry, port, promotion. They are often used as if they belonged to the same family of technical operations. They do not.

**The model registry is a versioned directory.** It records *what exists*: a model, its version, its training lineage, its validation metrics, its current status (none, staging, production, archived). It is

toolable: MLflow, DVC, Weights & Biases, SageMaker Model Registry, Vertex AI Model Registry provide industrial implementations. It is derived from the DevOps discipline and inherits its conventions: a table, versions, logged status transitions. Its function is useful, sometimes excellent. The problem is not its existence. The problem is its overestimation.

**The promotion port**, a concept I propose in this article as a doctrinal object, is an interface that contractualises *under what conditions what exists may cross a boundary*. It is not a status, it is an interface. It does not record the crossing, it *institutes* it. It answers a different question: not *what is in production?* but *what has been admitted, under what contract, with what revocation clause?* It contains entry criteria, validation methods, limits of validity, surveillance conditions, authorised evolution modalities, and withdrawal conditions.

**The promotion** is the event by which an artefact crosses the port. It is not a CI/CD operation, because a CI/CD operation transports code and weights, it automates movement. It is an act of institution, which recognises that a model, at a given moment and under named conditions, may be held responsible for the effects it will produce. The operation is technical; the promotion is institutional. A model may be deployed technically without being promoted institutionally. This is, in fact, the most frequent situation.

The distinction fits in one line. *The registry answers to what exists, the port answers to what may cross*. A directory does not make the boundary; a boundary without contract does not make the institution.

This distinction is not terminological. It conditions the ability to audit an AI system in regulated production. An auditor who arrives with the question *which model is running?* receives the answer from the registry. An auditor who arrives with the question *why this one, under what method, within what domain of validity, with what latitude for evolution, under what revocation condition?* needs the port. The second question is the one the regulator asks, not the first. And it is also the one many organisations do not yet know how to address.

### 3. Diagnosis: The MLOps Misunderstanding

The mainstream MLOps literature, from 2018 to the present day, has institutionalised the model registry and left the port unthought. Four factors, which this article gathers into an explanatory taxonomy, account for this asymmetry.

**First factor: the DevOps inheritance.** MLOps was built on the DevOps matrix and transposed its primitives: CI/CD, pipelines, staging and production environments. This transposition was necessary: it allowed AI to leave the artisanal notebook, a prior state in which governance often consisted of finding the right `final_model_v7_really_final.pkl` file. But it also imported a category error. In DevOps, code crosses the boundary through a branch merge and an automated deploy. The contract, if any, is implicit. Transposing this presupposition to AI amounts to asserting that a model which has passed its validation metrics is ready for production. This

is insufficient, because a model is not a service, it is an *estimator subject to the data distribution* it will encounter. Cross-validation on a test set is not a production contract, it is a snapshot of a past moment.

**Second factor: toolability.** The model registry prevailed because it was immediately codable. A registry is a table with rows and statuses. Platform teams delivered it within two quarters; data science teams adopted it because it replaced their Excel spreadsheets; executives validated it because they could understand it. The port is not codable without an architecture that carries it, because it presupposes defined boundaries, an admission contract, a revocation clause, responsibilities distributed across data science, business, quality, security, compliance, operations. It is not installable by package; it is *instituted* by a practice. What installs, propagates. What institutes demands power, time, and organisational courage.

**Third factor: the staging/production confusion.** In most stacks observable in 2026, models live in permanent staging, that is, in an intermediate zone where they run in parallel to production without ever being truly admitted, nor ever truly retired. The LangChain *State of Agent Engineering 2026* reports that 57% of organisations declare agents in production, but 32% cite quality as the primary obstacle to generalisation. What this means: many of these agents are in production in name, in staging in fact. The Deloitte *State of AI in the Enterprise 2026* formulates the same observation in another form: without a *hard gate* at 2x productivity lift, pilots become permanent. The grey zone is the degraded form of the port that was not instituted: it allows one to enjoy the effects of deployment without fully assuming the obligations of admission.

**Fourth factor: the cost of assessment.** Qualifying the admission of a model into regulated production presupposes an assessment that is not that of cross-validation. It mobilises prospective cohorts, silent trials, impact assessments, inter-cohort protocols, sometimes ethics reviews. This cost is out of reach of the CI/CD pipeline. It automates only partially, it demands an investment of expert time that a platform team cannot absorb without institutionalising it. The ease with which a model registry manages the *staging to production* transition masks precisely what instituting the port would cost. So long as the port is not a named object, its cost cannot be budgeted; so long as it is not budgeted, the platform team keeps crossing without contract. The registry is free, the port is expensive. This economic asymmetry explains the other asymmetry, that of tooling.

The consequence is observable. When the audit arrives, the organisation produces an extract of its model registry (such model, such version, such metrics). The regulator asks what has been opposed as admission contract. The organisation has nothing to produce. The registry attests to existence, not to admission. The difference is exactly the one that separates a directory from an institution.

## 4. Anatomy of the PCCP: The Contractualised Promotion Port

The Predetermined Change Control Plan, finalised by the FDA in December 2024 in its guidance *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions* (Docket FDA-2022-D-2628), constitutes today one of the most mature examples of what a contractualised promotion port should be.

It sits within a broader convergence: FDA/EMA principles on AI in drug development, entry into force of the QMSR aligned with ISO 13485, ramp-up of the European regime applicable to AI models. These regimes do not legally reframe the PCCP and do not make it enforceable by themselves. They nonetheless signal a shared functional requirement: to make explicit the conditions under which an AI system may evolve without breaking its regime of safety, quality, or responsibility.

Its object is precise: to allow a medical device manufacturer integrating AI-based software functions to describe in advance certain future modifications of the device, together with the method for developing, validating, and implementing them, so as to avoid requiring a full new submission for every planned and bounded change. The PCCP does not merely say *we want to be able to change the model*. It says *here are the planned changes, here is how they will be produced, here is how their impact will be assessed, here is why they remain compatible with the safety and effectiveness of the device*. It contractualises evolution before it occurs.

The PCCP consists of three elements that a model registry can at best keep as attachments, but which it neither produces nor institutes by construction.

**The first is the *description of planned modifications*.** The manufacturer does not submit a frozen model. It submits a model *and* the evolution trajectory it reserves the right to follow without new submission. This description is positive, bounded, and intelligible: it enumerates what may change (parameters, training data, recalibration modalities, functional scope, sub-populations, thresholds) and implicitly excludes what may not. This logic transforms the *version* into a *regime of versions*. The question changes: no longer *which version is approved?* but *which family of evolutions remains admitted without breaking the initial contract?*

**The second is the *methodology of development, validation, and implementation of the announced modifications*.** This is not an intention, it is a protocol. The manufacturer commits to *how* it will produce the modifications, not only to *which* ones. A change implemented outside the declared protocol is not covered, even if it falls within the category of authorised changes. Two models may display the same overall performance while resting on very different constitutive processes. The metric result does not suffice to establish institutional equivalence. The method is part of the contract.

**The third is the *impact assessment*,** the anticipated evaluation of the effects the modifications will produce on the safety and effectiveness of the device. The manufacturer is not merely authorised to change; it has documented, *ex ante*, why the change does not break the initial safety

contract. This assessment is the instrument that distinguishes the PCCP from a blank cheque: it introduces a logic of anticipated responsibility.

These three elements together achieve something that MLOps *staging/production* does not achieve: they contractualise the *authorised evolution before it occurs*. The regulator does not receive a list of models, it receives an admission contract with its evolution clause. The passage into production is no longer a technical act: it becomes the signature of an enforceable engineering contract.

The PCCP artefact is exactly what the promotion port must be. It is not a status in a registry, it is an interface that contractualises. It does not record a version, it institutes a regime of admissible versions. It does not answer to *what is in production?*, it answers to *under what contract is this held there?*. The PCCP is not the universal promotion port, it is a situated regulatory form proper to the medical device field. But its architectural type is generalisable. The PCCP is not a procedure, it is a theoretical artefact that the FDA built before the architects named it.

## 5. The Hexagonal Prerequisite and the Discipline of the Silent Trial

A port presupposes a boundary. A boundary presupposes an architecture that defines it. One does not lay a promotion port in an architecture where the business core, orchestration rules, and external dependencies are entangled, because no surface would then exist for the port to rest upon.

Hexagonal architecture, formulated by Alistair Cockburn in 2005 in his essay *Hexagonal Architecture* ([alistair.cockburn.us/hexagonal-architecture](http://alistair.cockburn.us/hexagonal-architecture)) and presented equivalently under the name *ports and adapters architecture*, defines precisely this surface. The application core (business rules, use cases) is separated from the outside (databases, services, UI, ML models) by ports, each implemented by substitutable adapters. Cockburn's original formulation states the structuring rule: *create loosely coupled components that can be connected to their environment by means of ports and adapters*, so that components are interchangeable at all levels and test automation is facilitated. The 2026 literature transposes this pattern to ML systems and data platforms. Implementations where the ML model is an adapter behind an application port can be found at Thoughtworks, at Dev3lop, in the Python tutorials devoted to ML pipelines.

The promotion port installs itself in direct continuity with this discipline. It is not a code pattern. It is the institutional consequence of an architecture that has defined its boundaries. Where hexagonal architecture guarantees that the model is *substitutable*, the promotion port guarantees that it is *admissible*. One governs the technical; the other governs the institutional. Without the first, the second is theatre; without the second, the first is an engineering exercise without governance.

The silent trial completes the apparatus. The clinical field has matured, over a decade, an emerging discipline of prospective non-interventional validation: deploy the model in real production, in passive mode, without its outputs affecting the clinical decision, over a defined period, in or-

der to capture its behaviour on field data without suffering the consequences. The scoping review published by *Nature Health* in 2025, which examined 891 articles between 2015 and 2025 and retained 75, documents this practice while insisting on its persistent heterogeneity (protocols vary, durations differ, success criteria are not fully stabilised, formal guidelines remain to be consolidated). One must therefore avoid presenting the silent trial as a mature and universal standard. It must be qualified correctly: an emerging discipline of prospective qualification, particularly suited to medical AI systems or systems assimilable to them, when a model must be observed on real data before interventional admission.

Its value lies in its position within the promotion port. It is not any kind of experimentation. It constitutes a contractualised window in which the candidate model is confronted with reality without yet obtaining the right to act. It transforms the production environment into a non-interventional qualification bench. It does not merely say *the model performed well on a historical dataset*, it says *the model maintains acceptable behaviour when it meets the living flow of the system it claims to serve*. Recent clinical cases (AI-based limb length measurement, PubMed 40990984; hospital decision systems) document sequences of shadow trial followed by clinical trial, at a cadence that has no formalised equivalent in enterprise MLOps.

This is precisely what a model registry cannot require by construction. A registry can preserve the result of a silent trial. It cannot, on its own, make the silent trial a condition of admission. The port can.

Two practical conditions therefore emerge:

**First condition:** an architecture that materialises the boundary between business core and AI adapters.

**Second condition:** a prospective discipline that qualifies the crossing before admission.

Pharmacovigilance, medicine, and medical devices have progressively institutionalised these two dimensions (because the consequences of a bad inference there cannot easily be dressed up as agile lessons learned). Mainstream MLOps has largely left them out of its standard tooling.

## 6. Objection Addressed: Is the PCCP Transposable Outside Health?

The objection is serious and deserves a frontal response. The PCCP, one will say, is a situated regulatory object. It belongs to the medical device field and draws much of its force from the existence of an external authority (here the FDA) that may receive, evaluate, accept, contest, reject the plan. Outside this frame, no one signs the contract *ex ante*, no one enforces it, and the model registry remains the only available form of collective memory of crossings. Generalising the PCCP lesson outside health would produce a seductive but inoperative analogy: without an authority that receives, no possible institution, only governance theatre.

The objection is sound. It avoids the usual slide which consists of taking a very situated regu-

latory artefact, emptying it of its legal force, then applying it everywhere with the subtlety of an administrative stamp. But it does not destroy the thesis. It forces its refinement. It is not a matter of copying the PCCP outside health, it is a matter of transposing its *type*.

Three answers can be opposed to this objection.

**First answer.** Ex ante institution does not require an external authority *that receives*; it requires a device *that binds*. Pharmacovigilance itself demonstrates this beyond the strict PCCP. The Pharmacovigilance System Master File (PSMF) is an internal document, not subject to prior external homologation, which binds the organisation on its pharmacovigilance system. The *Guideline on Good Pharmacovigilance Practices (GVP) Module II, Pharmacovigilance System Master File*, in its revision 2 published by the EMA, defines the PSMF as a single document continuously maintained, structured in seven modules (QPPV, holder's organisational structure, safety data sources, IT systems, pharmacovigilance processes, system performance, quality plan) whose explicit function is to document compliance with obligations arising from Regulation (EC) No 726/2004 and Directive 2001/83/EC. Documentation does not come from an authority; it is held by the Marketing Authorisation Holder and is enforceable against the authority upon inspection. The ALCOA++ principles (Attributable, Legible, Contemporaneous, Original, Accurate, Complete, Consistent, Enduring, Available) define the quality of the internal traces required. These devices contractualise without an external agency signing them. The authority, when it arrives, verifies whether the commitment has been kept. Institution precedes authority, it does not derive from it.

**Second answer.** The AI Act applicable to providers of general-purpose AI models does not reproduce the PCCP, and one must refrain from claiming otherwise. It rather imposes documentary and methodological pressure: technical documentation, transparency on training content, compliance policies, and, for models with systemic risk, evaluation and mitigation of systemic risks. The AI Transparency Atlas published in December 2025 (arXiv 2512.12443) evaluated fifty models on a weighted grid where safety disclosures count for 25% and critical risks for 20%; frontier laboratories peak at 80% compliance, the majority falls below 60%. The gap with an enforceable PCCP regime is documentable. The ramp-up of GPAI obligations will not invent the port; it will reveal which organisations have already instituted theirs and which have only their registry.

**Third answer.** The PCCP is not to be copied; its *type* is to be transposed. What it teaches is the structure of an artefact that contractualises: description of what may change, methodology for doing so, impact assessment, surveillance, withdrawal conditions. This structure is agnostic to the field of application, even if its criteria are sectoral. A credit scoring model in banking production can build its own artefact (description of authorised retraining, admissible variables, fairness tests on protected classes, drift thresholds, suspension conditions). A predictive maintenance model on critical infrastructure will describe the sensors covered, operating ranges, environmental conditions, extrapolation limits. A defensive model will describe autonomy level, prohibited action classes, human safeguards, deactivation mechanisms. The type is shared. The

content is situated. The organisation that builds it *ex ante*, when the authority arrives, will have an artefact to produce. The one that has not built it will have only a registry to produce, and the extract will not stand in for admission.

The objection turns against itself. To say the PCCP is intransposable amounts to saying that *ex ante* institution is a privilege of fields regulated from above. It is the exact opposite. Fields regulated from above built the institution because it was imposed on them. The others deferred. The deferral is coming due.

## 7. ToxTwin as an Implementation Terrain

ToxTwin, as a pharmacovigilance-typed digital twin, is exposed to the convergence of the three regimes enumerated in the introduction (FDA/EMA Guiding Principles of 14 January 2026, QMSR ISO 13485 in force since 2 February 2026, ramp-up of the European regime applicable to AI models between 2025 and 2027). It offers a terrain where the question of the promotion port is posed not abstractly but operationally. It stands at the crossroads of several disciplines (pharmacovigilance, AI, health data, software quality, traceability, auditability), which makes it precisely the type of system in which the absence of a port becomes visible.

The V2.4 architecture of ToxTwin laid the hexagonal frame. The application core is the pharmacovigilance business logic (adverse event reporting, causality, MedDRA classification, interface to individual case safety reports). The adapters cover ingestion of source data, linguistic decoding services, classification models, export interfaces to national regulatory systems (EudraVigilance, FAERS). An AI model (for example an imputability classifier or a signal detection engine) plugs in as an adapter behind an application port. This discipline is essential: if the model becomes the core, the organisation eventually adapts the business to the constraints of the model, which silently inverts the governance hierarchy.

The promotion port, within this architecture, is not an optional extension; it is the surface that distinguishes a *candidate* model from an *admitted* model. Its design imposes several constraints, which can be unfolded on a first sub-case: the imputability classifier.

It imposes the *explicit description of authorised retraining*: what data scope, what frequency, what learning window, what method for revising the ground truths. What is not in the description is not authorised, even if the pipeline technically makes it possible. The pipeline confers capability; the port confers the right.

It imposes the *methodology of prospective validation*: a silent trial of a duration calibrated on the volume of signals per week, with a requirement of inter-cohort stability before going live. This duration cannot be decided abstractly, because it depends on case volume, product variability, context criticality. A two-week silent trial may be derisory; a six-month silent trial may be disproportionate. The port contractualises duration against risk and volume, not against committee patience.

It imposes the *impact assessment on the target population*: expected effect on the true signal detection rate, on the false positive rate, on the distribution of imputability classes. A model that improves an overall metric at the cost of degradation on sensitive populations (elderly patients, polymedicated, rare effects) does not cross the port. The assessment said so *ex ante*.

It imposes, finally, a *revocation clause* (a fourth component I propose as an extension of the three-component FDA frame). The PCCP implicitly contains the logic of a bounded evolution regime; the promotion port must explicitly formulate the withdrawal regime. Revocation conditions include statistical drift, calibration degradation, major modification of the reference framework, critical pharmacovigilance event, regulatory change affecting model validity. A port without revocation is an admission without exit, which is not governance, it is lazy authorisation.

A second sub-case illustrates the generality of the treatment. A signal detection engine in pharmacovigilance (disproportionality algorithm on EudraVigilance data, Bayesian estimator of the *Information Component* type (Bate, Lindquist, Edwards, Olsson et al., Uppsala Monitoring Centre, 1998), weighted relative risk score) raises homothetic but distinct questions. The description of authorised evolutions concerns, here, the recalculation window, the retained disproportionality threshold, the list of MedDRA terms included in active surveillance, the rules for merging therapeutic classes. The methodology of prospective validation must contend with a phenomenon proper to signal detection: false positives are only observable after investigation, and investigation consumes expert time. The silent trial here is therefore not a passive shadow, it is, in the protocol I propose, a calibration window where the engine runs in parallel and where triggered investigations are conducted but not attributed to the pipeline in place. The impact assessment measures early detection lift, additional investigation load, confirmation rate after closure, distribution of new signals across therapeutic classes. The revocation clause covers at minimum two scenarios: performance drift relative to the baseline, and alteration of the EudraVigilance base itself between two extractions (an observed, not theoretical, case). Believing that regulatory data is a stable floor because it appears institutional is a costly naïvety.

These two sub-cases (imputability classifier and signal detection engine) are not redundant. They show that the promotion port is not a generic contract, it is a contract *situated by the use it qualifies*. A port is not decreed at platform level; it is declined at each application port it guards. This reinforces, rather than weakens, the thesis: the hexagonal surface is the only frame in which this multiplication remains governable, because it reserves to each application port its admission contract without imposing a generic contract at the core level.

This instance illustrates what the thesis claims. It does not prove it universally, it is not an argument, it is a case attesting that the construction is possible, anchored, and enforceable against a regulator arriving tomorrow. Other instances, in other fields, would produce other materials. The instance does not substitute for the argument.

## 8. Doctrinal Consequences

The promotion port modifies the way AI industrialisation must be thought. Five doctrinal consequences follow.

**First consequence: production is no longer an environment, it is a regime.** In many organisations, *production* still designates a technical environment (exposed endpoint, active monitoring, required availability, collected logs). This view is insufficient for regulated AI. Production is a *usage regime* in which an artefact produces enforceable effects. It is not enough to be deployed there; one must be *admitted* there.

**Second consequence: governance cannot be added after deployment.** An organisation discovering, at the moment of audit, that it must justify the admission of a model has already lost part of the battle. It can reconstruct justifications, aggregate evidence, produce notes; it cannot cleanly recreate the *ex ante* institution. Late governance resembles defensive archaeology. The promotion port compels documentation *before*, not after. It transforms admission into a condition of deployment, not a retrospective commentary.

**Third consequence: performance is only one clause of the contract.** A model may be performant and inadmissible. It may achieve a better AUC and remain uncalibrated. It may reduce mean error and worsen distribution tails. It may improve a global metric and produce an unsustainable operational load. It may work on a majority population and fail on critical cases. The promotion port forces one to leave the religion of the score, because it demands a situated judgement on usage.

**Fourth consequence: revocation becomes as important as admission.** A serious organisation does not only define how a model enters production. It defines how it leaves. Withdrawal must not be improvised in a moment of crisis; it must be foreseen as a normal clause of the admission regime. This discipline is still rare in mainstream MLOps. Technical rollback exists; institutional de-admission, far less. These are two different gestures: reverting to an earlier version is not withdrawing the right to act from a model.

**Fifth consequence: the promotion port is a management object, not only an architectural one.** It engages the CTO, the quality head, the business, compliance, security, sometimes legal, sometimes medical or regulatory. It forces the organisation to decide who may admit, who may block, who may revoke. It ends the fiction that production of a model is a purely technical matter. This fiction long served many parties well, by allowing the business not to understand, the data scientists not to govern, the platforms not to arbitrate, and executives to ask *where are we on AI?* the way one asks where one is on the network infrastructure.

## 9. Conclusion

The promotion port is a doctrinal object before being a technical one. Confusing it with the model registry is a category error that explains why so many organisations, in 2026, declare models in production without holding the contract that would make this an admission. The FDA built the artefact before the platform architects named it; it did so in a field regulated from above because it was compelled to, but the structure it produced is not specific to medical devices: it is the generic form of the institution of a crossing.

Two practical implications.

**First implication.** Organisations industrialising AI in regulated or regulation-candidate environments gain by building their promotion port *ex ante*, without waiting for the arrival of the authority. The PCCP provides the structure (description of authorised evolutions, methodology, impact assessment, revocation clause). Hexagonal architecture provides the surface. The silent trial provides the discipline of prospective qualification. The three together constitute what I propose to name the tripartite apparatus of the promotion port. None of the three, in isolation, is sufficient.

**Second implication.** The ramp-up of AI regimes between 2025 and 2027 will be a revealer. Organisations that have instituted their port will have an artefact to produce. Those that have only a registry will pull a table extract and discover that the auditor was asking for something else. The difference between a directory and an institution is not visible in peacetime; it is at the moment of audit.

A pipeline logs into a registry; a twin stands on a port. The difference is not one of tooling. It is one of regime.