From Certified Product to Continuous Authority

Why the PCCP and AI Health Regulation Document the Unfinished Emergence of a New Qualification Regime

The certified product, the classical paradigm of industrial qualification, presupposes a punctual act of authority: an object is compliant at a given instant, and that compliance holds until explicit revocation. Evolving AI puts an end to this paradigm. When a system can be modified without ceasing to be the same product, the punctual act of authorisation loses its bearing. This article documents the institutional emergence, partial and unfinished, of the paradigm that replaces it: continuous authority over an evolving system. The terrain where this mutation is observable is AI health regulation. Its central institutional moment is the Predetermined Change Control Plan (PCCP) finalised by the FDA on 4 December 2024.

I. Two Facts, a Sectoral Cluster, Three Generations: A Displacement to Qualify

Two Level 1 facts and a Level 2 sectoral cluster converge in April-May 2026 on AI health regulation.

1. **First element, Level 2.** The FDA, through its Center for Devices and Radiological Health, publicly maintains an *Artificial Intelligence and Machine Learning (AI/ML)-Enabled Medical Devices* list recording the devices authorised since 1995. This list is an official transparency instrument, not consolidated by the FDA itself into analytical statistics. Sectoral analyses aggregating the list's entries report, as of end-2025, a cumulative count of approximately 1,451 authorised devices, including 258 over 2025 alone, with a sectoral concentration on the order of 76% in radiology, the remainder distributed across cardiology, neurology, oncology, pathology, and an aggregate of other specialties. These consolidations fall under Level 2 in the present analysis's hierarchy of evidence: they rest on the FDA list, but their aggregation is sectoral.
2. **Second fact, Level 1.** On 4 December 2024, the FDA publishes the final guidance *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions*. For AI-enabled device software functions, the FDA formalises a mechanism that permits the bounded modification of an authorised system without a new complete submission, provided the modifications are described in a Predetermined Change

Control Plan reviewed by the FDA within the initial marketing submission. Three components are prescribed: the *Description of Modifications*, the *Modification Protocol* (covering development methods, verification, validation, and implementation), and the *Impact Assessment*. This is a hard fact sourced in the official guidance.

3. **Third fact, Level 1.** On 10 February 2025, the FDA addresses a warning letter to Exer Labs concerning unauthorised AI claims. This is a documented case of public enforcement action on an AI device, and it establishes that the FDA can act against a poorly qualified or poorly marketed AI product. Its doctrinal reach is limited and it should be presented for what it is: minimal evidence of post-market authority, not yet strong evidence of an accomplished regime of continuous authority.

Methodological note. The factual coordinates below are verified at primary FDA source for the PCCP guidance of 4 December 2024 and for the Exer Labs warning letter of 10 February 2025. Statistics relating to the aggregation of the FDA list, to PCCP adoption, and to AI-enabled recall rates derive from transparent secondary sources (academic publications, sectoral reports with declared methodology). Where an external coordinate is a Level 2 signal rather than a Level 1 hard fact, the text says so.

An author's analytical grid. To order the authorisations cumulated between 1995 and 2025, I propose a periodisation into three regulatory generations. This grid is analytical, not officially FDA. Its value is interpretive, and its refutability rests on the coherence of the inflections documented at each threshold.

1. **First generation, 1995-2015: software absorption.** The earliest medical algorithms are qualified as software, with no recognised AI specificity. The FDA handles the object under the preexisting 510(k), De Novo, and PMA regimes, without distinct evidentiary requirements.

2. **Second generation, 2016-2020: static ML identification.** The FDA progressively identifies machine learning as a distinct regulatory object, publishes in 2019 a framework discussion paper, but still treats the models as frozen post-authorisation. The logic remains: the model deposited is the model qualified, full stop.

3. **Third generation, 2021-2026: emergence of adaptive governance.** The *Predetermined Change Control Plan* concept appears, progressive publications of the SaMD action plan, finalisation of the PCCP guidance in December 2024. For AI-enabled device software functions, the FDA formalises a mechanism that permits the bounded modification of an authorised system without a new complete submission, provided the modifications are described in a Predetermined Change Control Plan reviewed by the FDA within the initial marketing submission.

A critical distinction before going further. Not all AI devices are evolving. A majority of the authorisations record *AI-enabled* systems that are static at the time of authorisation: a frozen model is filed, qualified, deployed. The institutional mutation described in this article does not concern the entire stock. It concerns the structuring minority of *adaptive AI* devices, or of static devices placed under a PCCP that authorises their future evolution. It is this minority that forces the displacement.

A conceptual precision. When I speak in what follows of *continuous authority*, the term is a convenient editorial label. The underlying concept designates the relationship between a regulator and an evolving system endowed with four infrastructural properties:

1. It is *temporally extended* (not punctual, prolonged beyond the initial act).
2. It is *reactive to observed deviations* (drift, incidents, version changes, context modifications).
3. It is *mutually binding* (the regulator can revoke or impose corrective actions; the operator must declare and maintain the evidence).
4. It is *evidentiarily reconstructible and opposable*, meaning that eight components must be documentable at any moment in the system's life: exact model version, dataset or evaluation distribution applied, timestamped inference logs, runtime configuration, operational clinical context, associated human decision where applicable, trace of PCCP modification with its three components where invoked, post-modification impact evidence.

At this stage, the PCCP partially institutes properties 1, 2, and 3. Property 4 remains a horizon of empirical validation. *Continuous authority* is therefore, as of 2026, an analytical category whose most advanced institutional incarnation is a proto-institution.

Three facts converge. The FDA has not accomplished the passage to the paradigm of continuous authority. It has instituted one of the first recognisable forms of it. What remains is to distinguish what is observable from what remains prescriptive.

II. Procedural Plurality, Evidentiary Reconfiguration, Specific AI Jurisdiction

The most common interpretive error in the public debate consists in conflating inherited procedural plurality with specific AI stratification. Three levels of analysis must be rigorously distinguished.

1. **First level, inherited procedural plurality.** The 510(k), De Novo, and PMA regimes preexist AI. They have structured all medical devices for decades: pacemakers, MRI machines, surgical scalpels, interventional devices. Their coexistence does

not, in itself, constitute an AI stratification. The FDA applies to AI submissions procedures it also applies to non-AI objects. If the demonstration stopped here, the classical sectoral objection would be sound: there would be nothing doctrinally specific, just ordinary medical regulation applied to new objects.

2. **Second level, internal reconfiguration of evidentiary criteria.** Within each of the preexisting procedures, have the evidentiary requirements for AI submissions been modified? Several FDA guidance documents and recommendations published since 2019 (Good Machine Learning Practice for Medical Device Development, SaMD action plan, PCCP guidance) progressively shift attention toward the description of training datasets, validation distribution, targeted operational conditions and, where a PCCP is attached, the anticipated trace of modification. These shifts do not all assert themselves as uniform general obligations across every procedure, but they progressively structure the expectations in ML-enabled submission practice. Internal reconfiguration does not consist in changing the procedure; it consists in changing what the procedure is invited to require as evidence.
3. **Third level, emergence of a specific AI jurisdiction.** Does the FDA institutionally recognise AI as a regulatory object with its own regime?

As of 2026, partially, and it must be disambiguated. Three sub-types of jurisdiction exist and must not be confounded.

1. *Sectoral jurisdiction augmented by AI*. A preexisting sectoral authority that adapts its requirements to the AI object, without ceasing to be sectoral. The current FDA falls into this sub-type for AI medical devices. The authority remains *FDA medical device*, and AI is an attribute of the regulated object.
2. *Horizontal AI jurisdiction*. An authority whose proper object is AI, independent of the sector of application. The European AI Act and the British (AISI UK) and US (CAISI) AI Safety Institutes fall into this sub-type. The proper object is the AI system as such.
3. *Composite sectoral-plus-horizontal jurisdiction*. A regime in which a sectoral authority and a horizontal authority articulate themselves on the same object. This is the emerging regime in Europe for AI medical devices: Medical Device Regulation (MDR) plus In Vitro Diagnostic Regulation (IVDR) plus AI Act, with sectoral notified bodies, EUDAMED, and composite vigilance. As of 2026, the operational mirroring of this composite regime remains partial. Europe is mobilised here as mirror, not as a second terrain of demonstration.

The sectoral concentration observed in the FDA portfolio (on the order of 76% in radiology according to the sectoral aggregations) suggests a complementary hypothesis. The internal evidentiary reconfiguration does not unfold uniformly across classes of use.

Computational radiology, technically and institutionally more mature, concentrates the most structured requirements. Other classes of use (CDS, generalist clinical decision support, continuous monitoring) remain in a zone of adjustment.

The sectoral-absorption counter-thesis, formulated as follows, is serious: *« The FDA has merely absorbed AI into its classical medical regime. The continuous authority doctrine is an ex post rationalisation of ordinary regulation. »* This objection does not fall at the first level (where it is well-founded), nor at the third (where the FDA does remain sectoral). It falls at the second level, and only at the second: the displacement of evidentiary expectations is sufficiently documented to speak of internal mutation, not pure absorption. The PCCP, which we now examine, is the most advanced demonstration of this mutation.

III. PCCP as Proto-Institution of the Mutation

The Predetermined Change Control Plan, in its verbatim title *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions*, is an institutional recognition, by a major regulator, that an evolving AI device is not a frozen product. It is a process whose authority must be maintained over time.

The mechanism is short to describe. The operator submits, alongside its initial authorisation request, a plan that bounds the future modifications it will be permitted to execute without a new complete submission. This plan comprises three components prescribed by the guidance of 4 December 2024.

The *Description of Modifications* enumerates the anticipated modifications, with their scope, their frequency, and the conditions under which they will be triggered. A retraining of the model on more recent data, the adjustment of a decision threshold, the extension to a new population distribution, the addition of an auxiliary function: everything must be described with sufficient precision so that an FDA evaluator can understand what the operator allows itself to do and what it forbids itself.

The *Modification Protocol* describes the methodology that will be applied to each modification. Development methods, verification procedures, required clinical or technical validation, implementation plan, traceability. This component is the most technical and probably the most structuring. Within the PCCP's perimeter, it institutes a *meta-procedure* governing future changes.

The *Impact Assessment* documents the analysis of the possible consequences of each modification: effect on clinical performance, on patient risk, on the overall quality of the system. This component institutes the doctrinal traceability of the plan: every modification carried out under PCCP must correspond to an impact documentation reusable in case of post-incident investigation.

What the PCCP institutes, by construction, is precise:

1. First, the recognition that system modification is anticipated and planned, not suffered.
2. Second, the bounding of the authorised variability within a pre-specified envelope previously evaluated by the FDA.
3. Third, the FDA's retention of the capacity to revoke the PCCP if the conditions cease to be satisfied.
4. Fourth, evidentiary continuity through every modification, via the *Impact Assessment*. The PCCP therefore partially satisfies properties 1, 2, and 3 of the continuous authority concept.

For property 4 (enforceable evidentiary reconstructibility, articulated to the eight components posed in §I), the PCCP sets out the conditions without guaranteeing their public effectiveness.

What the PCCP is not sufficient to institute also deserves to be stated. It is a bilateral regulator-operator instrument. It does not, as it stands, open a public procedure of appeal or third-party contestation. The six publicly contestable elements identified in Volume 4 (qualification perimeter, declared footprint, claimed comparability, maintenance of authority after incident, change of version, evaluator's conflict of interest) find in the PCCP a partial operationalisation for the fourth and fifth. The others remain to be articulated.

Industrial adoption of the PCCP must be measured before drawing doctrinal conclusions. The figures available to date are Level 2 signals, that is, from transparent academic and sectoral publications rather than from consolidated official FDA statistics. As of the third quarter of 2025, the reported cumulative count is approximately 67 PCCPs, of which 59 are recorded over the last two years. For the year 2024 alone, 41 PCCPs are reported to have been submitted within the ML-enabled 510(k) perimeter, that is approximately 1% of all 510(k) clearances of every kind, but 16.7% of clearances specifically ML-enabled that year. This double reading is doctrinally important. At the scale of the global medical-device portfolio, the PCCP remains highly minoritary. At the scale of the adaptive ML-enabled segment, it already exceeds one-sixth.

If adoption stagnates below 15% of the adaptive ML-enabled segment at a 36-month horizon, the mutation remains institutionally available but marginally operationalised. The thesis of partial emergence survives this scenario. It reformulates, without collapsing, as *« recognised but underused mutation »*. If adoption crosses 30% at the 36-month horizon, the mutation ceases to be marginal and becomes structurally deployed.

The PCCP is one of the first explicit institutional forms, within the perimeter of AI-enabled device software functions, by which a major regulator recognises that an evolving AI

system is a process whose authority must be maintained over time. It institutes three of the four properties. What remains is to render the fourth publicly contestable, and to extend its operational use.

IV. Five Strata and Graduated Post-Market Effectiveness

AI health regulation does not stratify a single thing. It stratifies five, which must be distinguished before examining which is doctrinally the most structuring.

1. *Usage stratification*. Imaging, CDS, continuous monitoring, computational histopathology, interventional robotics. Five typological classes observable, with radiology dominating according to sectoral aggregations.
2. *Evidentiary stratification*. Data requirements, clinical validation, drift instrumentation, dataset diversity, human supervision. Differentiated by class of use and by procedure.
3. *Surveillance stratification*. Post-Market Surveillance regime, drift management, recall protocols, real-world performance reporting. Differentiated by criticality.
4. *Insurance stratification*. Insurability, actuarial pricing, exclusions and coverage conditions. Differentiated by risk profile.
5. *Jurisdictional stratification*. FDA versus MDR-IVDR-AI Act versus PMDA versus Health Canada, with uneven mirroring.

The central argument of the present section is the following: the insurance stratum is beginning to become an analytical pillar, because it does not produce a normative discourse. It produces an enforceable price of risk. A price is an opposable economic decision. When an insurer requires *traceability*, *drift governance*, *human override*, *reconstructible monitoring* as conditions of insurability, stratification ceases to be a discourse and becomes an infrastructural economic constraint.

Three insurer positions document the insurance-side construction of AI risk, with health and HealthTech included.

1. Vouch identifies a structural gap in the HealthTech market: *« Most Cyber and Tech E&O policies in the insurance market explicitly exclude bodily and personal injury, signifying a gap in coverage that could leave HealthTech companies vulnerable. »* The position is doctrinally loaded: an insurer publicly declares that an entire swathe of AI health risk is not, as things stand, covered by existing policies. The economic consequence is direct: a HealthTech operator without coverage on bodily risk has no complete insurance qualification, whatever its clinical scores may be.

2. Coalition announces, conversely, an affirmative extension of the cyber perimeter to AI: *« The new endorsement expands the definition of a security failure or data breach to include an AI security event, where artificial intelligence technology caused a failure of computer systems' security. »* The extension is doctrinally important: cyber coverage now assimilates AI failures to classical security failures. Stratification is built no longer by documented exclusion but by typological inclusion.
3. W.R. Berkley documents the symmetric position by explicit exclusion: *« any actual or alleged use of AI, including any product or service sold by a company incorporating the technology »*. When an insurer explicitly excludes the use of AI, insurance qualification collapses for the operators concerned. Stratification operates through selection.

Four other documented insurer positions (Munich Re, Lloyd's framework 2025, plus two other sectoral sources) complete the picture. With seven documented insurer positions with verifiable clauses, exclusions, and endorsements, the insurance stratum becomes a structuring signal. The argument remains inductive: it proceeds from HealthTech, from cyber AI, and from generic E&O exclusions, toward AI health devices. It is not yet direct evidence of the actuarial pricing of AI medical devices, but evidence of the constitution of the AI insurance basin around exclusion, endorsement, and selection.

A strategic distinction is here imperative to avoid a frequent confusion. Three objects must be rigorously separated when speaking of AI health risk.

1. First, the risk of the AI health device itself: this is what the manufacturer incurs in product liability or in errors and omissions. This is the object of the insurance stratum discussed here.
2. Second, generic cyber and AI risk, outside health specificity: mobilised as a structuring analogy, not as a central case.
3. Third, the use of AI *by* the health insurer to deny or approve claims, of which the *Lokken v. UnitedHealthcare* case and the PXDX controversy at Cigna are the most visible illustrations: this third object does not fall under the insurance stratum, but under the jurisdictional stratum, because it poses a problem of procedure and public contestability, not of insurability. Conflating it with the first two would blur the doctrinal argument.

The post-market effectiveness of FDA regulation can be examined through a graduated six-level scale, which usefully replaces the single nuclear test of revocation.

1. **Level 1, documented enforcement actions.** The Exer Labs case, warning letter of 10 February 2025 for unauthorised AI claims, illustrates the FDA's capacity to intervene publicly against a poorly qualified or poorly marketed AI product. It

establishes minimal post-market authority evidence on this category of stakes. It does not, alone, prove the maturity of a sophisticated regime of algorithmic surveillance.

2. **Level 2, recalls linked to software or algorithm.** A study published in *JAMA Health Forum* reports that 79 recalls, that is 43.4% of the recalls identified over the studied period, occurred within the 12 months following initial clearance. The *recall-free survival rates* reported are 96.6% at one year post-clearance, 93.5% at three years, and 91.8% at five years. Publicly listed companies are involved in 90% of documented events. Source to be reconfirmed at primary reference at the time of final drafting, Level 2.
3. **Level 3, post-market corrective actions, restrictions on claims, reinforced reporting obligations.** These actions are exercised case by case and call for a systematic empirical census. To be documented for final publication.
4. **Levels 4, 5, and 6** (PCCP adoption above 15% of the adaptive segment, post-incident revocation, public operationalisation of property 4) constitute the higher indicators of the scale, whose materialisation at a 3-5 year horizon would decide the degree of accomplishment of continuous authority.

An underlying economy deserves to be named. Sustaining continuous qualification requires specific infrastructural capacities: dedicated safety teams, runtime instrumentation, reconstructible PMS capability, opposable evidentiary archiving. These capacities are costly, and their cost is cumulative. The capital threshold of continuous qualification becomes a structural barrier to entry, analogous to the one described in the previous volume for benchmark-breakable qualification. Observable stratification is not only typological. It is capitalistic.

V. Emerging Tension, Terminal Limit Signalled

The current FDA regulatory architecture is device-centric. It qualifies a circumscribed object, operating on a defined clinical domain, with localised instrumentation. The PCCP does not change this logic. It extends it in time, without changing the regulatory unit.

But emerging clinical AI systems are composed. A generalist foundation model supplies the cognitive core. A clinical wrapper specialises it. An orchestration coordinates several sub-systems. The whole has no obvious regulatory localisation.

Three consequences. The cognitive core is neither stable nor locally governed. A modification at the upstream operator (foundation model update, prompting system change, guardrail adjustment) may affect clinical behaviour at the downstream operator without FDA submission. The clinical wrapper alone may be qualified, but its qualification

does not capture the variability of the underlying foundation model. The PCCP, designed for a single device, must be doctrinally extended to handle a chain of operators.

This tension is signalled as a terminal limit, not developed as analytical core. It will be the object of a distinct volume, because it poses a proper conceptual question: what becomes of the *promotion port* of an AI artefact when that artefact is composed, and when its operators are disjoined?

Three graduated falsifiability conditions close the present section. If at a 36-month horizon no effective modification is documented under PCCP, the mechanism remains theoretical. If adoption stagnates below 15% of the adaptive ML-enabled segment, the mutation remains marginal. If enforceable evidentiary reconstructibility is never operationalised by a public procedure, continuous authority remains an analytical category without complete incarnation. The present thesis survives the first two conditions in modalised form. The third demands substantial reformulation.

VI. The AI Governance Pentalogy and the Historical Rupture

Five volumes now compose the AI governance pentalogy. Physical constraints. Differentiated allocation. Artefact promotion. Qualification as infrastructure of concentration. Partial emergence of the maintenance of continuous authority. The sequence sketches a progressive displacement of the AI governance question: from material to temporal, from resource to relation, from certified to maintained.

The present volume documents the moment when this displacement becomes observable in a critical sector. AI health regulation is not the future of AI governance. It is, for the evolving segment, its partial present state. Nor is it an automatically transposable model. No sector today combines with as much visibility a public stock of authorised AI devices, a predetermined change mechanism, and sectoral post-market surveillance. But the trajectory is sufficiently documented to be examined without being prophesised.

For COMEX members and CTOs of large health groups, the operational consequence is precise. Not to debate the adoption of a continuous qualification regime. To internalise it. This presupposes reconstructible PMS capacities, operational drift instrumentation, documented version governance, and opposable evidentiary archiving infrastructure. For regulators of other critical sectors (regulated finance, autonomised aviation, defence AI, energy), the consequence differs. Anticipate. Prepare the infrastructures, competencies, and capital required for continuous qualification before sectoral regulators come to require its opposability.

For insurers and actuaries, the moment is doctrinally decisive. Seven documented insurer positions demonstrate that the market is beginning to transform AI risk into a

price. The enforceable price is one of the mechanisms through which stratification becomes real. Not principled discourse. Not regulator-to-regulator rhetorical mirroring. The price. The extension to specific product liability or E&O contracts dedicated to AI health devices remains the next step to observe.

The history of critical human infrastructures presents, without determinism, a recognisable regularity. Pharmacovigilance took three decades to institute itself as a continuous regime after thalidomide. Continuous authority over evolving AI systems could follow an analogous trajectory, accelerated or slowed by political, capitalistic, and sectoral conditions. Health documents the mutation. It does not guarantee it.

Health does not yet demonstrate accomplished continuous authority. It shows something more interesting: the moment when a regulator, confronted with evolving systems, begins to transform authorisation into relationship. The PCCP is not the end of this mutation. It is the first recognisable institutional form of it.

Volume 5 of the AI Governance Pentalogy. Previous volumes: "Energy as a Governance Constraint" (vol. 1), "Allocating the AI Kilowatt-Hour" (vol. 2), "The MCP Vulnerability as a Pure Case of the Promotion Port" (vol. 3), "When Evaluation Becomes the Infrastructure of Concentration" (vol. 4).