FOUNDATION

1

Challenge of the Ages

Of all the transformations a CxO has led, this one is different.

1.1 Transformation Like No Other

On April 2, 2026, the U.S. Food and Drug Administration (FDA) issued the first warning letter in its drug CGMP (current good manufacturing practice) history with a section titled “Inappropriate Use of Artificial Intelligence.” The firm, a homeopathic over-the-counter drug manufacturer, had used AI agents to draft drug specifications, procedures, and master production records, and had distributed product without conducting required process validation. Asked why, the firm told investigators the AI agent had not informed them the requirement existed. The agency rejected the posture, cited U.S. 21 CFR 211.22(c), and made clear that any AI output used in CGMP activities must be reviewed and cleared by an authorized human representative of the firm's quality unit. The company has since ceased drug production.

This is the first warning letter of its kind. It will not be the last. The architectural commitments every regulated enterprise has to make from this point forward must answer the question that quality unit presumably could not: when AI produced this, who in this firm knew enough to recognize that it was wrong?

In a 2025 essay, Anthropic CEO Dario Amodei wrote a sentence that has stayed with me longer than most things I have read about artificial intelligence.

> *“Humanity is about to be handed almost unimaginable power, and it is deeply unclear whether our social, political, and technological systems possess the maturity to wield it.”*
>
> DARIO AMODEI · ANTHROPIC CEO · 2025

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

Most industries can answer Amodei's implicit question about readiness with a decisive and honest *not yet*. We cannot. For those of us inside regulated industries, that sentence lands differently. Our answer has to be *yes*.

We have spent careers in industries whose entire premise is that powerful tools, applied without maturity, hurt people. The artifacts differ (drugs and devices in my world, complex instruments in finance, airframes in aviation, reactors in nuclear), but the architectural answer converges. We do not invent the maturity to wield these things alongside their invention. We build it, slowly and painfully, in the architecture of how we work: design controls and validated systems, audit trails and ongoing surveillance, change control, escalation paths, named accountable humans. A regulated enterprise is, more than anything else, a system for not letting powerful things hurt people without our understanding why.

Now we are being handed something more powerful than anything in our history, and it does not arrive as a molecule, an aircraft, or a financial instrument. It arrives as *cognition*. It thinks alongside us. It writes our procedures and drafts our reports. It classifies the exceptions and suggests the responses. It is everywhere in the workflow at once, and every place it appears is a place where a single decision lands on a person who did not make it.

Every executive over forty has lived through several transformations. The internet. The mobile interface. The data lake. The shift from on-premise to SaaS and Cloud. Each felt disorienting in its moment. Viewed in hindsight, each was a *substrate* change: a shift in where the work runs, how it connects, what it costs, and who maintains it, but not a shift in the work itself. A clinical trial in 2010, run on paper case report forms and fax, is recognizable as the same activity as a clinical trial in 2020 run on a cloud-hosted electronic data capture system. Only the floor under it changed.

AI IS NOT A SUBSTRATE CHANGE

Artificial intelligence is not a substrate change. It is a *cognition* change, and that is a different category of event entirely. The floor under the work was already shifting in ways we had governance for. The work itself is now shifting in ways we do not.

When you put cognition into a workflow, three things happen at once that no prior transformation produced.

- **The locus of decisions moves.** Choices that previously sat in the hands of a particular reviewer (the medical monitor reviewing an adverse event, the quality assurance (QA) reviewer dispositioning a deviation, the

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

regulatory writer composing a submission narrative) begin to sit *partially* with the system. Not because anyone designed it that way, but because the system has now done most of the thinking before the human looks. The reviewer's role shifts from *deciding* to *agreeing or not agreeing*, and that is a different cognitive job, with different failure modes. That quality unit, presumably asked to certify outputs they had never been the principal author of, made that kind of mistake.

- **The speed of decisions accelerates beyond the speed of governance.** A traditional quality management system (QMS) was designed around decisions made by humans at human pace. A deviation reviewed in days. A corrective action closed in weeks. A change control completed in months. AI systems, left to themselves, will produce dispositions in seconds. If governance can only react in days, the mismatch creates a structural gap into which anything can fall. The fix is not slowing the AI. The fix is making governance run at the same speed.
- **Accountability dissolves into the network.** With a traditional system, when a decision goes wrong, you can trace the fault. You can identify the reviewer, the standard operating procedure (SOP) they followed or failed to follow, the training they had or lacked, the data they could or couldn't access. With an AI-mediated workflow, that traceability evaporates unless you have built it back in deliberately. Whose decision was the missed adverse event signal? The model's? The reviewer who agreed with the model? The data scientist who fine-tuned the model? The vendor who supplied the foundation? The QMS owner who validated the system? Without architecture, the answer is *no one*, and *no one* is not an acceptable answer. Regulators have been clear: it must be a named, qualified, accountable human, every time, and the architecture has to enforce it.

This is why the phrase *human in the loop*, useful as it has been, is not enough. A loop describes someone watching from the side, intervening when something looks wrong. That metaphor was sufficient for systems in which the human did most of the work and the machine assisted. It is insufficient for systems in which the machine does most of the work and the human reviews. The geometry has inverted, and the language has not caught up.

Read together, the regulations describe a different commitment, one this book reads as placing the human not in the loop but in the *core* (human-centered). Structurally, not procedurally. To give this commitment a name for the purposes of this book, we will call the operational form it takes the *Round Table*: humans and agents seated as peers in a single forum, contributing through defined roles, sharing the same context and lineage. Accountability for outcomes

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

is attributed to a named human at every consequential step. This is not a slogan. It is what the regulatory texts have already committed to, in plain language, and much of this book is the work of showing what that commitment actually requires.

1.2 Who This Book Is For

This book is for the leader of a regulated enterprise whose career depends on getting AI right. The title varies with organization size. Chief Executive Officer at a small life-sciences or fintech firm. Chief Compliance Officer, Chief Risk Officer, or Chief Digital Officer at the mid-size. Senior Vice President of AI Governance, or Chief AI Officer, at the largest. The position is constant: a leader who will be held accountable for the architectural commitments their organization runs under, and who can delegate downstream once they hold the analysis whole.

It is also a book for the consultants, architects, governance leads, and regulatory affairs officers who serve those leaders. The CxO commits the enterprise; the consultant advises and helps build. The ratio is something like one executive to several hundred advisors, integrators, and implementation partners whose work shapes what actually gets built. The book serves both populations: substantive enough to satisfy the practitioner, plainly written enough that an executive can hold the entire picture and ask the questions that drive their organization.

What this book is matters as much as *who* it is for. It is an analysis. Not a regulations reference, not a narrative business book, not a technical manual. It is a dissection of what the major regulators and standards bodies set out when they are read together. Regulators and standards bodies across three continents, and the leading advisory firms alongside them, describe the same architectural commitments for AI in regulated enterprise. The book's contribution is to show those commitments, category by category, with the joints between them and the inspection paths an auditor will traverse.

The analysis is deliberately multi-industry. It reads the regulatory influences of several heavily regulated sectors at once, and the inclusion of automotive and nuclear, among others, is not incidental. Many of the examples lean toward life sciences, because that is the world I have principally served. The value, though, is the convergence itself, and the projection of those shared commitments onto the architectural layers any regulated enterprise has to build. The layers are the same whatever the sector the examples come from.

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

Five sectors anchor the analysis, each grounded in its own instruments:

- **Life sciences** — U.S. FDA, the International Council for Harmonisation (ICH), and ISO 13485:2016 (medical-device quality management systems)
- **Automotive** — UN Regulation No. 156 (the UNECE software update management regulation)
- **Aviation** — the EASA AI Concept Paper, Issue 02 (the European Union Aviation Safety Agency's machine-learning guidance)
- **Nuclear** — IAEA SSG-39 (the International Atomic Energy Agency's safety guide for instrumentation and control systems)
- **Finance** — Federal Reserve SR 26-2 (the 2026 revised U.S. supervisory guidance on model risk management, superseding SR 11-7)

One signal is worth marking at the outset. When the U.S. banking agencies revised their model risk management guidance in 2026, they expressly placed generative and agentic AI outside its scope, leaving those systems to each organization's broader governance and flagging a further consultation to come. The newest guidance defers the AI-specific question to the architectural commitments an enterprise builds across its lifecycle, which is exactly the ground this book's analysis categorizes.

The thesis is structural. Organizations that navigate the AI era successfully, avoiding warning letters and patient harm, converge on the same architectural commitments because the regulations have already specified them. The book lays the analysis out so the categorization is in hand before the inspection arrives.

The promise is concrete. By the end of the book, you will hold the categorization the analysis lays out, and the obligations the regulations specify under each. You will recognize, when a vendor presents their AI platform, whether the answer they give to "how does this satisfy EU AI Act Article 14" is real or rehearsed. You will recognize, when a consultant scopes a transformation program, whether the categories of obligation they organize their work around match the categories an inspection will traverse. You will recognize, when a colleague reports that "the AI flagged this exception," whether the work behind that statement can withstand the same question a regulator now knows to ask.

You will not become a data scientist. You will not need to. You will become the kind of leader who can lead from understanding, ask the questions that matter, and recognize whether the answers hold. That is the promise.

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

The book delivers on it in three movements. The first three chapters establish the foundation: what AI's category-of-event shift means for a regulated enterprise, how the major regulators have already converged on the architectural commitments that follow, and how this analysis groups those commitments into five categories. The next five chapters walk the categories (data, thread, intelligence, team, and process), naming the responsibilities, the controls, the failure modes the regulators have already cited, and the convergent witness the advisory firms provide. The final two chapters are crosscutting. Chapter 9 lifts the categories that have assembled, and treats mutual validation, lifecycle governance, and cybersecurity posture as obligations that run across every layer. Chapter 10 closes by addressing the reader's first 365 days inside the analysis, with reading paths by role and a calendar from Day Zero forward.

1.3 The Stakes and the Imperative

The opportunity is real. McKinsey's life sciences analysis suggests that *up to 95% of life sciences roles may have agentic teammates*, and that *two-thirds of those roles will involve building, managing, and supervising agents.* That is not a small reorganization. It is the second-largest workforce reshaping in the industry's history, after the move to global outsourcing of clinical and manufacturing operations. The pattern repeats, with different artifacts, in every regulated vertical. The companies that get this right will run their pipelines faster, write their submissions cleaner, monitor their post-market signals earlier, and operate their plants with fewer deviations. They will also, and this matters, be better places to work. The routine cognitive load that exhausts scientists and clinicians will move to systems built for it, and the humans will be freed to do what only humans can do.

The risk is also real, and Gartner has named it sharply: *over 40% of agentic AI projects will be cancelled by the end of 2027, due to escalating costs, unclear business value, or inadequate risk controls.* The same forecast that gives us the opportunity tells us that nearly half of the attempts will fail. Some of those failures will be benign. Abandoned pilots, sunk-cost write-offs, embarrassed sponsors. Some will not. In a regulated industry, an inadequately controlled AI project does not just fail commercially. It produces warning letters, triggers consent decrees, causes recalls, loses trial sites. In the worst cases, it harms the patient at the end of the chain. And in this way, perhaps the projects that don't get cancelled could do the most harm.

There is a middle path between the 95% opportunity and the 40% failure rate, and it is not a strategy of caution. The companies that will win this transformation are *not* the ones that move slowly. They are the ones that move quickly *with*

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

the architectural commitments the regulations describe already in hand. They treat the regulator's expectations not as a brake but as a *blueprint*. They recognize that the architecture which satisfies the U.S. FDA and the EU AI Act simultaneously is the same architecture that produces the highest-quality outputs, the cleanest audit trails, the most efficient post-market surveillance, and ultimately the most trusted product. Compliance and capability are not in tension here. They are the same thing, viewed from two angles.

The companies that will lose are the ones that buy the platform first and ask the architectural questions afterwards. They will discover, eighteen months in, that the grounding their AI outputs need was never citation-bound, the override events that needed first-class capture were never recorded, the adaptation envelopes the regulations require were never authored, and the validator-agents that were supposed to qualify the work were never qualified themselves. They will be told by their auditors that the entire deployment is unfit for regulated use. They will have to choose between rebuilding from underneath or removing the AI from the workflow it has now become indispensable to. Neither is a happy quarter.

The window for getting this right is not a decade. The major regulatory milestones are now.

- **Quality Management System Regulation (QMSR)** — United States; effective February 2, 2026, the amended and retitled U.S. 21 CFR Part 820, now incorporating ISO 13485:2016 by reference.
- **EU AI Act** — European Union; high-risk obligations currently applicable August 2, 2026, with a provisionally agreed amendment (the Digital Omnibus on AI) expected to defer them to December 2, 2027 (Annex III standalone) and August 2, 2028 (Annex I embedded in regulated products); transparency obligations from August 2, 2026.
- **Predetermined Change Control Plan (PCCP)** — United States, FDA guidance; final guidance already issued.
- **U.S. FDA AI/ML Lifecycle** — United States, FDA draft guidance; expected to finalize in the Center for Devices and Radiological Health (CDRH) FY2026 queue.
- **ISO/IEC 42001:2023** — international standard, *Information technology — Artificial intelligence — Management system*; already published.
- **prEN 18286** — draft European standard; finalizing through 2026.

Every regulated vertical has its own parallel calendar; the urgency is the same. By the end of this calendar year, the regulatory architecture will be substan-

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

tially complete, and the question of whether your enterprise is built to run inside it will be answerable in a single inspection.

To meet that, an executive needs a particular kind of understanding. Not the vague sense that AI matters, the hype has done that work, and the board, the competition, and the trade press have done theirs. What is needed now is the substance: the specific, named, defensible understanding of what AI in your enterprise actually has to look like for it to be safe, compliant, effective, and worth the investment.

The substance is architectural. It is not a strategy deck, not a maturity model, not a vendor selection rubric. It is the structural shape that emerges when the regulations are read together, category by category, with each category's responsibilities, each category's failure modes, and each category's connection to the texts that already govern your industry. The architecture is what the building has to sit inside. The leader has to know it.

The leader who understands the design intent is the one who can tell when the implementation matches it. The implementation can be delegated. The architectural literacy that makes the implementation defensible cannot. When the inspector asks how this AI output was reviewed, when the board asks how the investment is governed, when the journalist asks whose decision the harm came from, the answer has to come from a person who understands the architecture, not from a deck someone else wrote.

The systems we need are coming into focus. Regulators across continents, advisory firms, and platform vendors describe the same architectural commitments in their own vocabularies. What remains is for an executive to hold the analysis whole and make their enterprise run inside the commitments it surfaces. This does not require us to slow down.

It does require us to lead.

> *"AI is a cognition change, not a substrate change. The executive cannot delegate cognition to a system the executive does not understand."*
>
> THE CENTRAL CLAIM OF CHAPTER 1

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

IF YOU REMEMBER NOTHING ELSE

AI is a cognition change, not a substrate change. *The locus of decisions moves, the speed of decisions accelerates beyond the speed of governance, and accountability dissolves into the network unless the system builds it back in deliberately.* The executive who understands the analysis leads through it. The executive who does not, is led by it.

NEXT · CHAPTER 2 — *Five regulators on three continents describe the same architectural commitments; the leading advisory firms have arrived independently at the same picture. The next chapter looks at the convergence.*

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org

FOUNDATION

2

The Regulators Have Spoken

The regulators have written what the architecture has to satisfy.

2.1 Regulators Are Clear

There is a moment in every transformation when the executive picks up the regulations directly and reads them alongside everything else. It is usually a clarifying moment. The regulators have been working on the AI question for some time, and their writing has the specificity, plainness, and durability that we'd hope for. When you read the AI guidance from the U.S. FDA's Center for Devices and Radiological Health alongside the EU AI Act, the European Medicines Agency (EMA) Reflection Paper, the International Medical Device Regulators Forum (IMDRF) N88, and the U.S. FDA, Health Canada, and the UK's Medicines and Healthcare products Regulatory Agency (MHRA) Good Machine Learning Practice (GMLP) guiding principles, you do not encounter five different futures. You encounter one future, described five times.

Consider the most consequential single sentence in any of these documents, Article 14 of the EU AI Act, the *human oversight* article. It requires that high-risk AI systems be designed and developed in such a way that they *"can be effectively overseen by natural persons during the period in which they are in use."* Now, even more important than the phrase itself, are the words *designed and developed* earlier in the sentence. The Article does not say that human oversight must occur. It says that the system must be designed for human oversight. It is an architectural requirement, not a procedural one. You cannot retrofit it by writing an SOP.

The Article goes further. It specifies what that design must enable. The natural person must be able to understand the capacities and limitations of the system, to remain aware of automation bias, to correctly interpret the system's output, to decide not to use the system in any particular situation, to override or reverse a decision the system has made, and, in the most architecturally specific

Purchased from RAPS Library (library.raps.org) for the exclusive use of unknown.
© 2026 RAPS Library. Please report unauthorized use to raps@raps.org