Verit Global Labs
10291 Sweet Bay CT
Parkland Fl, 33076

(954)871-7524
hero@veritglobal.com
www.veritglobal.com

March 9, 2026

BIOTECH & PHARMA AI STRATEGY

Building a Governed, Evidence-Driven Biopharma

AI is not a tools initiative — it is an operating model transformation. The companies that win will be those that build a governed, evidence-producing enterprise where every AI-supported decision is traceable, defensible, and inspection-ready.

Biopharma's bottleneck is evidence integrity at scale, not scientific capability.

"The ability to move a decision, a data artifact, a safety signal, or a regulatory submission… with traceability and defensibility" determines speed, compliance, and competitiveness.

AI strengthens—not replaces—scientific, clinical, and regulatory judgment by making it faster, more consistent.

Why AI Fails in Pharma Today

AI fails in pharma not because of models, but because of fragmented data: identity fragmentation, ontology divergence, version opacity, and provenance gaps.

"Deploying AI without solving the data foundation problem builds confident inconsistency at scale." The maturity model resolves this progressively, culminating in a unified data fabric.

This framework directly solves the core structural problem holding biopharma back — fragmented evidence, inconsistent decisions, and ungoverned AI. It gives a biotech the blueprint to transform into a governed, high-velocity, inspection-ready enterprise capable of scaling AI safely and competitively.

BIOTECH & PHARMA AI STRATEGY

Building an Enterprise That Discovers Faster, Executes with Evidence, and Scales Under Governance

An Executive Value Narrative · February 2026

The Central Thesis

The defining constraint in a modern biotech or pharmaceutical enterprise is not scientific talent, capital, or technology. It is evidence integrity at scale — the ability to move a decision, a data artifact, a safety signal, or a regulatory submission across the full value chain with enough speed, traceability, and defensibility to satisfy the most demanding audience that will ever examine it: a health authority inspector.

AI does not change what a pharmaceutical company must do. It changes what the organization is capable of doing with the people, systems, and institutional knowledge it already has. Done well, AI transforms every function across the value chain — from target identification through post-market pharmacovigilance, from clinical operations through commercial execution and medical affairs — into a more capable version of itself. Not by replacing scientific or regulatory judgment, but by making that judgment faster, better-supported, and consistently applied across the enterprise.

> ***The question is not whether to deploy AI in drug development. The question is whether to deploy it as a collection of unvalidated point tools — or as a governed, evidence-producing operating system for the full value chain.***

This narrative describes how to build the second outcome: a structured, eight-level maturity progression in which each level unlocks capabilities not safely achievable before and creates the conditions for the level that follows. Every domain is addressed — discovery, clinical, regulatory, manufacturing, quality, pharmacovigilance, commercial, medical affairs, market access, supply chain, and corporate functions — because the value chain is indivisible.

What Makes Pharmaceutical AI Uniquely Difficult

Every AI deployment decision must be understood against the regulatory architecture of the pharmaceutical value chain. The consequences of getting it wrong are not measured in user dissatisfaction — they are measured in patient safety, health authority enforcement, and the integrity of the drug approval record.

GxP Data Integrity Is Non-Negotiable

Every AI action that touches a GxP system — a clinical trial EDC, a QMS deviation record, a PV safety database, an eCTD submission artifact — must meet the same data integrity standards as any other computerized system: ALCOA+ attribution, contemporaneous recording, original data preservation, accuracy, and completeness. An AI agent that writes to a GxP system without a complete, replayable audit trail does not save time. It creates an inspection finding. The enterprise must treat agent actions as computerized system operations subject to 21 CFR Part 11 and EU Annex 11, not as convenient shortcuts around validation requirements.

Autonomy Carries Patient Safety Implications

In pharma, AI autonomy is also about patient safety and the validity of the clinical and manufacturing record. An agentic system that executes the wrong randomization, submits an unsupported regulatory claim, or misroutes a serious adverse event does not create a business problem — it creates a potential patient harm and an inspectable record entry that may never be fully correctable. Autonomy must be granted through a documented, evidence-based tier structure, not assumed from model capability claims.

The Evidence Chain Must Be Reconstructable End-to-End

When a health authority questions a decision made during a clinical trial, manufacturing deviation, or safety evaluation, the institution must reconstruct exactly what information was available, what analysis was performed, who decided, under what authority, and what the outcome was. AI operating outside a governed evidence chain creates a gap in the reconstructable record that no retrospective documentation effort can reliably close.

Commercial and Medical Functions Carry Their Own Compliance Burden

Commercial, medical affairs, and market access are not lighter-touch AI zones. Promotional compliance, fair balance, off-label communication, and consent management carry obligations that OPDP and health authorities actively monitor. An AI agent that suggests off-label content, generates promotional materials missing required fair balance, or creates field briefings violating territory compliance boundaries creates enforcement exposure as consequential as any GxP finding. These functions require their own governance track — running in parallel to the GxP track from Level 0.

The Data Foundation Problem: Why AI Alone Is Not Enough

The most underestimated obstacle to high-value AI in pharma is not model capability or governance policy. It is data — specifically, the fragmented, inconsistent, poorly governed data infrastructure that underlies most pharmaceutical enterprises. AI amplifies whatever data quality it is given. In a fragmented environment, AI scales the inconsistency.

Four Layers of the Problem

Identity fragmentation: the same compound, patient, site, investigator, supplier, or HCP exists as multiple inconsistent records across systems. The SIS has one identifier, the safety database has another, the clinical trial platform has a third. AI drawing on these sources simultaneously without a reconciled identity layer produces unreliable output by construction.

Ontology divergence: different functions use different controlled vocabularies and coding systems for the same underlying concepts. Clinical uses MedDRA. Manufacturing uses internal deviation taxonomy. Commercial uses therapeutic area classifications that do not map to clinical endpoint definitions. Without a unified ontology layer, AI operating across functions produces answers that are internally consistent within each silo but contradictory across them — and the contradictions are invisible until they reach a prediction model or cross-function decision.

Version opacity: controlled documents, study protocols, regulatory submissions, and scientific positions change over time, but change history is poorly maintained. AI without version awareness conflates current and superseded content, producing outputs that cite outdated positions as authoritative.

Provenance gaps: the lineage from a data point to its source — which study, which analysis, which database state, at which moment — is broken in most pharma data architectures. AI cannot produce defensible output if it cannot cite exactly where every fact came from.

How the Maturity Progression Solves It — Level by Level

Level 2: Evidence Memory establishes version-pinned, citable document repositories and a first-generation enterprise ontology. AI retrieves only from controlled sources. Provenance is captured on every retrieval. Version opacity and the first layer of ontology divergence are directly addressed here.

Level 3: Orchestration creates governed data lineage for every decision — which data, at which system state, under which policy version. Cross-system data flows are tracked and traceable rather than assumed.

Level 5: Agent memory architecture is explicitly governed: short-term task state, long-term case memory, and the knowledge RAG layer all operate under defined retention rules, PII and trial subject personal data masking, and access controls. Context loss between agent invocations — a primary source of inconsistent agent outputs — is eliminated by design.

Level 6: Continuous modeling surfaces identity and ontology conflicts as anomalies rather than letting them compound silently. Inconsistencies that manual operations would never surface appear as prediction errors that the AI program must investigate and route to ontology owners.

Level 7: The data foundation problem is fully resolved through a unified data fabric: a single authoritative identity layer across all systems, a governed enterprise ontology mapping terminology across GxP, commercial, and scientific domains, a self-updating scientific knowledge graph, and continuous data quality monitoring that treats ontology drift as an operational risk metric with defined alert thresholds.

> ***An enterprise that deploys AI without solving the data foundation problem is not building intelligence — it is building confident inconsistency at scale. The maturity model is also a data maturity model.***

This is why Level 2 is the most under-invested and most consequential level in the progression. Organizations that rush past it to deploy Level 5 agents or Level 6 predictive models consistently encounter the same failure: impressive pilots that degrade in production because the data foundation was fragile. The data foundation problem cannot be solved retroactively at Level 7. It must be addressed beginning at Level 0 — governance of which data AI systems may access — and built deliberately through every level that follows.

Two Distinct Dimensions: Maturity Levels and Autonomy Tiers

The framework operates on two dimensions that are easy to conflate but must be kept clearly separate. Confusing them is one of the most common governance errors in pharmaceutical AI programs.

Maturity Levels (0 through 7) — What the Enterprise Has Built

The eight maturity levels describe the organizational and technical infrastructure the enterprise has constructed: governance architecture, knowledge foundations, orchestration backbone, operating surfaces, and analytical capabilities. A Level 3 enterprise has built orchestrated, closed-loop workflows with traceable evidence. A Level 6 enterprise has built continuous predictive intelligence across the value chain. These are organizational states that require investment, change management, and demonstrated performance at each stage to advance.

Autonomy Tiers (T1 through T5) — What an Individual Agent Is Permitted to Do

The five autonomy tiers describe what a specific AI agent or workflow is authorized to do, independent of what the underlying model is technically capable of. A T1 agent reads and synthesizes but takes no system actions whatsoever. A T4 agent executes material workflow steps with segregation of duties, dual controls, and an immutable audit trail (Bounded Execution — Material). Tier assignment is made per agent, per use case, per deployment context, and documented in the agent's context-of-use statement.

The relationship between the two dimensions: higher maturity levels make it safe and operationally justified to deploy more agents at higher tiers. But having a high maturity level does not mean every agent should operate at high tiers. Each agent earns its tier through documented performance history against the metrics defined in the AgentOps framework, and every tier escalation is a formal change control event.

The Five Autonomy Tiers

Tier	Name	What It Can Do	Pharma Examples	Required Controls	Risk If Skipped
T1	**Advisory**	Reads, synthesizes, classifies. Zero system writes of any kind.	Literature synthesis; feasibility review; signal briefing; protocol gap analysis	Model logging; source provenance; human sign-off on outputs	Unverified AI content entering regulated decisions silently
T2	**Draft & Recommend**	Prepares structured drafts and recommendations. Human must approve and execute every action.	MVR drafts; PV narrative scaffolds; deviation descriptions; regulatory Q&A; MLR pre-screen; territory briefs; value dossier drafts	Human approval gate before any execution; versioning; immutable transcript; PII and personal data redaction in logs	Unreviewed drafts entering the GxP or promotional record
T3	**Bounded Execution — Low Risk**	Executes reversible, low-risk actions within explicit constraints. Sandbox-first before production.	eTMF gap task creation; site follow-up scheduling; EDC query dispatch; CRM task creation; MLR routing; contract compliance alerts	HITL during ramp; tool allow-list enforced at runtime; per-action throttles; automatic rollback; post-action audit record	Untracked system changes; orphaned tasks without human visibility

Tier	Name	What It Can Do	Pharma Examples	Required Controls	Risk If Skipped
T4	**Bounded Execution — Material**	Executes consequential workflow steps with segregation of duties and dual controls.	PV follow-up dispatch (approved templates); deviation routing to CAPA; RIM updates; labeling task routing; specialty pharmacy coordination	Dual approval; SoD enforced by architecture; validated tools; immutable audit trail; kill switch; policy-as-code caps	GxP record changes without defensible evidence trail; inspection findings on computerized system changes
T5	**Semi-Autonomous (Narrow)**	End-to-end loop within a narrow, well-instrumented process. Humans supervise exceptions only.	Non-GxP ITSM resolution; supply chain exception proposals; internal KB maintenance; low-risk commercial operations	Continuous drift monitoring; safety triggers; incident playbooks; re-validation on schedule; documented COU; regular access reviews	Silent scope creep into GxP workflows; performance drift undetected over time

Tier Escalation Is Change Control

Every escalation from one tier to the next requires documented performance evidence from the current tier, a risk assessment for the proposed tier, updated context-of-use documentation, and QA approval. An inspector who sees a well-documented tier escalation record is satisfied. One who sees an undocumented autonomy increase finds an undocumented computerized system change — which is a data integrity finding regardless of whether the agent performed correctly.

The Eight-Level Maturity Progression

Each level represents a categorical shift in organizational capability — not merely more AI features, but a different operating posture for the enterprise. The levels are cumulative: orchestration at Level 3 depends on the knowledge foundations of Level 2; autonomous agents at Level 5 require the governance architecture of Level 0 and operating surfaces of Level 4; predictive intelligence at Level 6 requires the operational data generated at Levels 3 through 5; and the unified data fabric of Level 7 requires the ontology groundwork begun at Level 2.

Level	Name	Core Capability	Signature Outcome
L0	Governed Foundation	GxP-aligned AI governance: COU framework, audit trails, policy-as-code, and promotional compliance governance across all functions	Shadow AI eliminated; every agent action reconstructable; GxP and promotional compliance posture established before scale
L1	Augmented Expertise	AI drafting and summarization embedded in expert workflows across every function — ELN, CTMS, safety DB, QMS, regulatory workspace, CRM, MLR platforms	Expert hours recovered from assembly work; documentation quality improves; draft-to-approval cycle compresses across all functions
L2	Evidence Memory	Governed, versioned, citable knowledge layer: controlled docs, submissions, claims, scientific positions, and first-generation enterprise ontology	Knowledge is authoritative, version-pinned, and citable; ontology inconsistencies surfaced; provenance captured on every AI retrieval
L3	Orchestrated Value Chain	Consequential decisions become governed cases with traceable authority and cross-system execution from IND through NDA and lifecycle	Manual handoffs collapse; evidence travels with every decision; exceptions are visible and owned; audit trail is continuous
L4	Role-Based Operating Layer	Unified, policy-enforced work surface for every function — CRA, PV processor, reg writer, quality associate, MSL, commercial ops, market access specialist	System-hopping eliminated; institutional knowledge embedded in every role's surface; policy enforced consistently; onboarding accelerates
L5	Bounded Autonomous Agents	AI agents execute scoped tasks within tier-governed policy caps — every action audit-ready, every tier assignment backed by performance evidence	40–60% cycle time compression on high-volume case work; expert capacity multiplied; AgentOps telemetry proves performance continuously
L6	Predictive Value Chain	Risk, quality, signal, and commercial data modeled continuously; ontology conflicts surface as prediction anomalies; leaders act proactively	Safety signals surface earlier; manufacturing drift detected before OOS; territory under-performance predicted; inspection gaps close proactively
L7	Adaptive Enterprise	Unified data fabric, governed enterprise ontology fully resolved, self-updating scientific knowledge graph, and continuous AI governance	Value chain operates as governed, self-improving system; data ontology problem fully resolved; structural competitive advantage compounds annually

What Each Level Delivers Across Every Domain

The matrix below shows what each maturity level delivers for each major enterprise function. Reading across a row shows the progression of capability for a single domain. Reading down a column shows what a given level enables simultaneously across all functions — illustrating why each level creates compound, enterprise-wide value rather than point-solution value.

How to read: Identify your current or target level and read down that column to understand simultaneous unlocks across all functions. Identify a domain of interest and read across to see its full maturity arc.

Domain	L0	L1	L2	L3	L4	L5	L6	L7
Discovery & Preclinical	Governed AI in ELN & lab tools	Literature synthesis; assay drafts; report scaffolds	Versioned scientific record; compound ontology	IND package orchestration; CRO workflow cases	Scientist copilot: protocol + evidence unified	Autonomous pipeline execution; lab agents	Candidate attrition prediction; target scoring	Self-updating scientific knowledge graph
Clinical Operations	GCP boundary defined	MVR drafts; query letters; ICF support	Current protocol always accessible; site ontology	Site activation & eTMF orchestrated end-to-end	CRA unified surface: CTMS + eTMF + queries	eTMF monitoring agents; data cleaning agents	Enrollment risk predicted 4-8 wks early	Integrated trial intelligence IND to CSR
Pharmacovigilance	GVP boundary defined	Case narrative drafts; MedDRA coding suggestions	Coding ontology versioned; prior cases citable	Intake routing & follow-up to GVP SLAs	PV processor workspace unified	Triage agents (T3/T4); follow-up dispatch	Signal emergence modeled continuously	Continuous post-market safety intelligence
Regulatory Affairs	Submission AI boundary defined	Q&A drafts; dossier gap lists; label notes	Prior submissions versioned; claims citable	HA correspondence cases; dependency tracking	Reg writer unified surface: dossier + draft	Submission readiness agents; RIM checks	HA intelligence synthesis; filing risk model	Continuous regulatory lifecycle management
Manufacturing & Quality	GMP boundary (Part 11/Annex 11)	Deviation drafts; CAPA scaffolds; SOP pre-checks	SOPs & batch specs versioned; one source	Change control & deviation case orchestration	QA unified surface: deviation + CAPA + CC	Batch review agents; EM trend monitoring	Manufacturing drift detected before OOS	Continuous quality governance; process AI
Supply Chain	Supply & SOX boundary defined	Shortage scenarios; S&OP pack assembly	Supplier specs & contracts versioned	Cold chain & chargeback case management	Planner surface: forecasts + exceptions	Supply exception agents within policy caps	Demand & supply risk predicted ahead of market	End-to-end supply intelligence; dynamic alloc
Commercial Operations	Promotional compliance AI defined	MLR pre-screen; territory briefs; call prep	Approved claims library; label language current	Label change impact on promo assets orchestrated	Sales surface: CRM + compliant content unified	MLR routing agents; field inquiry triage	Territory under-performance predicted early	Continuous commercial intelligence & compliance
Medical Affairs	Medical AI boundary; off-label controls	Congress synthesis; MI drafts; MSL pre-call briefs	Approved scientific content versioned & citable	Insight workflow; publication planning cases	MSL unified surface: KOL + content + notes	Insight synthesis agents; KOL mapping	KOL influence trends; evidence gap detection	Continuous scientific landscape intelligence
Market Access	Market access AI boundary defined	Value dossier drafts; tender response assembly	HEOR evidence library versioned; formulary data	Payer account cases; contract compliance tracking	Access team surface: formulary + HEOR + contract	Contract compliance agents; access barrier detection	Formulary shift & access patterns predicted	Outcome-based contract evidence platform

Key insight: Every level creates value in every domain simultaneously. An investment in Level 2 (Evidence Memory) gives version-pinned knowledge to clinical, regulatory, quality, PV, commercial, medical affairs, and market access at the same time. This is why the infrastructure levels — 0, 2, 3 — have disproportionately high compound return relative to their direct cost.

Level 0: The Governed Foundation

> ***In pharma, governance is not overhead — it is the precondition for every unit of AI value. An agent that cannot be fully reconstructed and defended at inspection is a liability regardless of how accurate it is.***

The most consequential failure mode in pharma AI is not a model that produces poor output. It is an AI system that operates without a context-of-use definition, without audit trails, without change control, and without a clinical safety owner — and then surfaces in an FDA or EMA inspection as an undocumented computerized system. Internal risk assessments consistently find that 60–80% of AI-related regulatory risk comes not from the AI systems themselves, but from the absence of governance, documentation, and oversight.

LEVEL 0 The Governed Foundation

GxP-aligned governance, audit architecture, promotional compliance controls, and policy-as-code before deployment scales

GxP Governance Architecture

- AI Steering Committee with named clinical safety owner: cross-functional authority (Regulatory, Quality, Medical, Legal, IT, Privacy) with explicit decision rights, risk classification authority, and escalation paths for GxP and commercial compliance separately
- Context of Use (COU) framework: every AI system assigned intended use, autonomy tier (T1–T5), risk class, GxP boundary, validation requirements, and named system owner before any deployment
- Risk-tiered intake: patient-safety-critical, GxP material, bounded execution, and advisory-only tiers each with distinct control packs; promotional compliance and off-label tiers for commercial and medical functions governed in parallel
- Policy-as-code enforcement: allow/deny lists, action caps, segregation of duties, and approval gates enforced at runtime by the policy engine — not by convention or manual discipline that erodes under volume pressure

Promotional Compliance Governance (Parallel Track)

- Promotional compliance lead named alongside clinical safety owner: accountable for OPDP boundary definition, MLR AI governance, and off-label communication controls from day one
- Approved content controls technically enforced: AI systems in commercial and medical functions may only retrieve, reference, and recommend content from the approved claims library and labeled indications — enforced at the policy engine, not by user discipline
- Off-label detection: conservative classifier active on all AI outputs in commercial and medical workflows; ambiguous cases escalated to medical information; no AI-generated response to off-label inquiries without medical review
- MLR AI governance: AI pre-screening of promotional content is governed as a distinct computerized system with its own COU, validation artifacts, and audit trail separate from the human MLR approval record
- Field communications archive: every AI-assisted

What Skipping This Costs

- Shadow AI in GxP workflows creates undocumented computerized system dependencies that surface as Part 11 findings or data integrity observations — typically discovered during inspection when remediation is most expensive
- Promotional AI without governance creates off-label exposure, fair balance violations, and OPDP enforcement risk that can disrupt commercial operations and trigger Warning Letters
- Ungoverned vendor AI creates data residency, personal data and trial subject data handling, and model-training-on-your-data risks that become legal liabilities and audit findings in both GxP and commercial contexts
- Retrofitting governance after incidents consistently costs 3–5x more than building it correctly upfront and carries reputational risk that no subsequent investment can fully recover

▪ Immutable audit logging: replayable transcripts of every agent action, tool call, input, output, and rationale with tamper-evident storage meeting ALCOA+ requirements across all functions ▪ Change control pathway for AI: model, prompt, and policy updates promoted through a controlled pipeline with QA approval and tested rollback capability — applied identically to GxP and promotional compliance workflows ▪ Vendor risk management: privacy posture, training-on-your-data risk, data residency, and contractual protections assessed before any model is deployed in any function	field communication — emails, call notes, MSL interaction records — archived with the AI system identifier, output version, and human approval record	

Governance Is the Fast Path, Not the Slow One

Organizations that invest in Level 0 governance before scaling consistently deploy faster at Levels 1 through 5 because the intake process is standard, the validation path is defined, and the escalation authority is clear. Those that skip it spend the same time — but after an incident, under regulatory scrutiny, at a far higher remediation cost.

Level 1: Augmented Expertise Across Every Function

> ***The fastest path to visible, credible AI value in pharma is to compress the assembly and documentation work that consumes expert time — across every function simultaneously, without changing validated workflows or creating inspection exposure.***

Scientific and regulatory expertise is the scarcest resource in a pharmaceutical enterprise. Across every function the pattern is the same: experts spend a significant fraction of their time on structured assembly work — gathering data that already exists in validated systems, formatting it against templates that already exist, and producing documentation whose structure is already defined. CRAs draft MVRs from CTMS data they have already reviewed. PV processors structure narratives from safety database fields they have already assessed. Regulatory writers compose HA responses from prior submissions and study data the scientific record already holds. Medical science liaisons assemble pre-call briefs from CRM history, approved slide libraries, and published papers accessed through separate tools. None of this is low-value work — it requires domain knowledge and judgment — but a substantial portion of it is assembly that AI can perform faster and more consistently.

LEVEL 1 Augmented Expertise

AI drafting and summarization inside existing validated workflows — expert quality maintained, assembly time compressed, audit trail preserved

GxP Functions (all T2 — Draft & Recommend)

- ELN experiment note summarization and protocol drafting with SOP constraints, reagent pre-population, and prior assay library reference — scientist reviews and approves before any ELN entry; AI system identifier recorded in the lab record
- Clinical monitoring visit report drafting from CTMS notes, site metrics, and prior findings — CRA reviews, edits, and approves; no auto-upload to eTMF under any tier
- PV case narrative drafting from structured safety database fields — processor reviews, edits, and electronically signs; every AI draft labeled as AI-assisted in the safety database record
- Regulatory Q&A drafting with cited evidence from prior submissions and study data — RA team reviews before any submission; claim-checking required on

Commercial, Medical & Market Access (all T2)

- MLR pre-screen: AI flags high-risk language, missing references, and incomplete fair balance in promotional assets before formal MLR submission — reduces reviewer workload by filtering content the brand team can correct before submission
- Territory and account briefing pack compilation: non-promotional account insights, call planning context, and approved content from CRM data — privacy and compliance filters active on every output
- Medical information response drafting with citations to approved label and published data — Medical Affairs reviews before sending; off-label escalation enforced automatically; approved sources only
- Congress insight synthesis: key scientific themes, cited abstracts, and strategic action items — medical

Why This Is the Trust-Builder

- Rework rate — the fraction of AI-assisted drafts requiring major revision before approval — is the leading quality indicator and the evidence base for every tier escalation above this level; measure it from day one
- When experts experience AI as giving them time back and producing work they can actually use, they become active participants in the broader program rather than passive observers or silent resisters
- The cultural dividend of Level 1 is as important as the efficiency dividend: functional teams that trust AI at Level 1 adopt Level 4 operating surfaces and Level 5 agents with substantially less friction
- Deploying Level 1 across all functions simultaneously — GxP and commercial in parallel — establishes the governance maturity and adoption baseline the enterprise needs before investing in Levels 2 and 3

every draft; citation to source document mandatory
- Deviation and CAPA description standardization from QMS inputs with root cause library scaffolding — QA reviews and approves before record is finalized; no auto-issuance under any circumstances
- CSR section population from controlled TLFs, protocol, and SAP outputs — medical writing team reviews and finalizes; trace map from data to narrative generated automatically and retained with the document

director reviews before any strategic use; commercial insights kept strictly separate from medical output
- Value dossier and tender response drafting from approved evidence library, clinical data summaries, and HEOR models — health economics team reviews before any payer submission
- S&OP meeting preparation: agenda, exception summary, and scenario options auto-compiled from planning systems — leadership reviews and decides; decision log attributed to human approvers with mandatory reason codes

Level 1 Use-Case Portfolio

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Target landscape synthesis	Monitors literature continuously; extracts claims; builds a traceable evidence graph of ranked target hypotheses	**T2**	Mandatory citations; IP boundary restrictions; scientific review before hypothesis enters formal target list	Coverage completeness; time from publication to synthesis; hypothesis-with-evidence rate
Assay design assistant	Drafts assay protocol with SOP constraints, reagent pre-population, and prior assay library check	**T2**	Template locking; lab head approval required; unit and parameter validation vs. instrument constraints	First-pass success rate; rework rate; time from request to approved protocol
Preclinical report drafting	Assembles study report from locked data extracts with full trace map from data to narrative	**T2**	Locked data pulls only; GLP controls if regulated; study lead sign-off required	Authoring time; audit finding rate; template compliance rate
MVR drafting (ClinOps)	Drafts monitoring visit report from CTMS notes, site metrics, and prior findings	**T2**	Template locking; CRA sign-off required; no auto-upload to eTMF	Drafting time reduction; rework rate; issue capture rate vs. manual
PV case narrative drafting	Drafts narrative from structured safety database fields with mandatory source grounding	**T2**	Processor sign-off required; 'no new facts' guard; locked pulls from safety DB only	Authoring time; quality score; rework rate
Regulatory Q&A drafting	Drafts response with cited evidence from prior submissions; claim-checking active on every draft	**T2**	SME review required; citation mandatory; no auto-submission under any tier	Response cycle time; acceptance rate; citation audit pass rate

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Deviation/CAPA drafting	Standardizes deviation description; scaffolds root cause from library; proposes corrective actions	T2	QA approval required; no auto-issuance; severity rubric enforced	Cycle time; classification accuracy; recurrence rate on closed records
MLR pre-screen	Flags high-risk language, missing references, and incomplete fair balance before formal MLR	T2	Conservative thresholds; no AI clearance of content; human MLR authority preserved; audit trail on all flags	Review cycle time reduction; rejection rate at formal MLR; compliance findings on cleared assets
Medical information response drafting	Drafts balanced response with citations to approved label and published literature	T2	Approved MI database only; off-label escalation enforced; medical review required before sending	Response time; off-label escalation rate; quality scores
Territory briefing pack assembly	Compiles account status, formulary position, prior call notes, and approved content for pre-call review	T2	Compliant content filter active; personal data and trial subject data redaction; no prohibited information inclusion	Preparation time; adoption rate; call quality scores
Value dossier drafting	Assembles HTA submission from clinical data, RWE, and HEOR models; routes for HEOR team review	T2	Controlled evidence sources only; HEOR team review; no unvalidated economic models	Preparation time; HTA submission acceptance rate; evidence request fulfillment speed
Cold chain excursion investigation draft	Classifies excursion severity from stability data; drafts investigation package for QA review	T2	QP disposition authority preserved; quality approval required before disposition	Time to disposition decision; investigation completeness rate
MSL field note synthesis	Extracts structured insights from MSL interaction reports; tags by therapeutic area, sentiment, competitive signal, and safety observation; assembles for medical director review	T2	Source-specific de-identification where required; medical director review before any strategic use; no commercial use of medical engagement data	Insight extraction time; actionable insight rate; medical director adoption
Publication planning and manuscript drafting support	Drafts structured publication outlines, tracks author coordination milestones, maps deadline risks, and scaffolds manuscript sections from study data and prior publications	T2	Full author control enforced; no undisclosed AI authorship; ICMJE guidelines active; human author sign-off required on all content before submission	Outline-to-draft cycle time; author coordination overhead; deadline adherence rate

Level 2: Evidence Memory and the Ontology Foundation

> ***Most regulatory and quality failures in pharma begin as a knowledge problem. Someone referenced an outdated SOP, cited superseded data, used a non-current template, or could not find the prior submission containing the answer. Level 2 makes knowledge authoritative, versioned, and consistently accessible — and begins to resolve the ontology divergence that makes cross-function AI unreliable.***

The knowledge infrastructure of a large pharmaceutical enterprise is sprawling and fragile. SOPs exist in multiple versions across functions and sites. Protocol templates diverge study-by-study. Regulatory submissions reference scientific positions that have evolved without a systematic record of how and why. The terminology used in clinical differs from the taxonomy used in manufacturing differs from the vocabulary used in commercial. AI operating across these fragmented sources amplifies their inconsistency at the speed of inference.

Level 2 builds the evidence memory layer: a governed, versioned, citable repository of controlled documents, scientific positions, prior submissions, and institutional knowledge — and a first-generation enterprise ontology that begins to resolve the terminology fragmentation that makes cross-function AI unreliable. This is a read-only layer that informs work; it does not execute it. Its value compounds at every level above it, and it is the single most important investment the enterprise can make to improve Level 6 model accuracy before any predictive model is built.

LEVEL 2 Evidence Memory

Governed, versioned, citable knowledge with version pinning, provenance tracking, and first-generation ontology management

Controlled Document Memory

- Version pinning enforced at retrieval: AI retrieves only approved, current document versions — never draft or superseded content; source provenance (document ID, version number, retrieval timestamp) captured on every retrieval and retained with the AI-assisted output
- Document ownership and freshness monitoring: stale, unowned, or conflicting documents flagged and routed to named owners before they produce operational errors or inspection findings
- Cross-reference integrity maintained systematically: SOP-to-SOP and protocol-to-SOP cross-references tracked in the system rather than by individual staff memory; updated automatically when source documents are revised

Scientific and Regulatory Memory

- Prior submission and scientific position library: regulatory responses and label claims anchored to the submission of record with full citation tracking and change history — answerable within minutes during HA Q&A cycles
- Evidence package repository: study reports, integrated summaries, and analysis datasets versioned and cross-referenced to the current dossier state — accessible to regulatory writers without manual reconstruction
- HEOR evidence library: health economic models, real-world evidence studies, and outcomes data versioned and citable for market access submissions with methodology documentation preserved

First-Generation Ontology Management

- Terminology mapping across functions: a governed vocabulary layer maps clinical MedDRA terms to commercial therapeutic area classifications, manufacturing deviation taxonomy to quality system codes, and regulatory indication language to commercial promotional claims — eliminating the cross-function contradictions that fragmented vocabularies produce
- Coding system governance: MedDRA version pinning for PV; ICD version tracking for clinical; ATC codes for market access; indication-specific terminology for commercial — each locked to the approved version for each workflow context;

▪ Approved claims library for commercial and medical: cleared promotional and scientific claims, each with source document, approval date, and expiry trigger — the knowledge foundation for compliant MLR pre-screening and governed AI recommendations at Levels 4 and 5	▪ Knowledge gap analytics: what questions cannot be answered, which document areas are underrepresented, where institutional knowledge exists only in individual memory — driving a systematic quarterly improvement program	version changes governed as change control events ▪ Entity registry: compounds, indications, sites, investigators, and HCPs maintained as a governed identity reference that all AI systems retrieve from rather than inferring from system-specific identifiers — the foundation of the unified identity layer completed at Level 7 ▪ Ontology gap detection and routing: when AI cannot reconcile a term across domains, the gap is logged, classified by risk, and routed to the appropriate knowledge owner; the ontology improves with every gap resolved, compounding accuracy for every AI system above it

Data Foundation Connection: Level 2 is where the data ontology problem is first actively addressed. The terminology mapping and entity registry established here are the foundation of the unified data fabric completed at Level 7. Organizations that invest seriously in Level 2 ontology governance find Level 6 predictive models significantly more accurate — because cross-function signals resolve to the same underlying concepts rather than colliding as unrecognized synonyms.

Level 3: Orchestrating the Value Chain

> ***The dominant failure mode in pharmaceutical operations is not poor decisions — it is poor follow-through. Site activation items untracked. PV follow-ups undispatched. Deviations closed without matching CAPAs. Label changes whose impact was never mapped to downstream commercial assets. Level 3 closes these loops by design.***

The value chain is a sequence of high-stakes handoffs: from discovery to IND, from protocol to site activation, from site to eTMF, from safety case to PSUR, from label change to global material update, from contract signature to compliance monitoring. At each handoff, information must travel between systems, evidence must be preserved, and the next action must be owned and tracked. When managed manually — through email, shared spreadsheets, and informal coordination — these handoffs fail at a predictable rate. In GxP contexts, those failures are inspection-reportable.

Level 3 converts every consequential handoff into a governed, closed-loop case: the action is tracked, evidence travels with the decision, exceptions are escalated with full context to the appropriate owner, and the system of record is updated once from a single approved action with a complete audit trail. This is orchestration with embedded governance — not RPA executing tasks in isolation, but a decision-and-evidence architecture that produces reconstructable records at every step.

LEVEL 3 Orchestrated Value Chain

Consequential decisions become governed cases — evidence travels, exceptions are owned, systems update once with full audit trail

Clinical & Safety Orchestration

- Site activation: missing document chasers, checklist tracking, and cross-party coordination through CTMS integration — every action timestamped and attributed; activation not marked complete without evidence of every checklist requirement
- eTMF completeness management: continuous scanning for gaps and misfiled documents; corrective task routing and resolution tracking to completeness thresholds — no agent-initiated TMF filings under any tier
- PV case routing: ICSRs triaged, de-duplicated, classified by seriousness and expectedness, and routed to correct queues with SLA tracking against GVP reporting clocks — T3 tier with full audit trail on every routing decision

Regulatory, Quality & Commercial Orchestration

- Submission readiness checks: pre-validation of packages for completeness and cross-reference integrity, gap routing before submission deadlines — right-first-time rate tracked by submission type and submission team
- Labeling change impact orchestration: every label change traced to all downstream assets across all markets — local labels, PIL, promotional materials, patient support materials — with task creation for each affected asset and completion tracking to regulatory timeline requirements
- Health authority correspondence: incoming queries classified, routed, and commitment tracking initiated with deadline monitoring and escalation for approaching due dates —

The Audit and Evidence Architecture

- Every orchestrated case carries a complete decision record at the moment of action: who acted, under what authority, with what evidence, referencing which policy version — captured structurally, not assembled retrospectively
- Exception queues surface operational risk in real time through dashboards — replacing the status meetings and report cycles that currently approximate this visibility, days too late
- Reconstruction for inspection is immediate: what happened, when, who decided, what evidence was available — answerable within minutes from the system of record, not from an archaeological email search
- Segregation of duties enforced through the case workflow architecture

- PV follow-up orchestration: follow-up questions drafted from case gaps, dispatched via approved templates after processor approval, responses tracked to GVP timelines — no case closes with an outstanding follow-up
- EDC data query management: discrepancies detected, queries drafted with context, batched for DM review, and follow-through tracked to resolution with contribution to database lock metrics

every commitment documented and closed

- Deviation and CAPA case management: standardized description, root cause scaffolded, CAPA tasks tracked to closure, recurrence monitoring engaged — no case closed without evidence that the root cause is addressed
- Change control orchestration: change requests mapped to impacted processes, documents, and validation states — every dependency tracked and dual-approved before implementation trigger
- Contract compliance monitoring: commercial contract term adherence checked continuously across active payer agreements; anomalies routed to market access and finance for human resolution; SOX controls active on all financial exception handling

rather than relying on individual discipline that erodes under volume pressure, staff changes, and peak workload periods

- Cross-function visibility without manual aggregation: one exception-focused view of active risk replaces the siloed status reports that miss cross-function dependencies until they become deadline crises

Level 3 Use-Case Portfolio

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Protocol feasibility agent	Matches protocol to site and patient population reality; ranks sites with evidence; flags enrollment risk	**T2**	Bias audits; data freshness checks; human approval of final shortlist	Time to shortlist; actual vs. predicted enrollment; screen fail rate
Site activation orchestrator	Tracks all checklist items across regulatory, legal, and operational workstreams; coordinates parties via governed comms	**T3**	SoD; approved communication templates; CTMS audit trail; no automated regulatory submissions	Activation cycle time; checklist completeness at activation; inspection-ready packages
eTMF completeness agent	Continuously scans eTMF against expected document inventory; drafts corrective task notifications; tracks resolution	**T3**	No auto-filing; controlled vocabulary for classification; sampling QC on AI classifications	eTMF completeness %; gap resolution time; inspection finding rate on TMF
SAE intake and routing	Classifies SAEs for seriousness and expectedness; routes to correct queue; drafts follow-up questions with medical context	**T3 / T2**	SLA rules; escalation for high-severity; human medical assessment required	On-time SAE reporting rate; follow-up completeness; time to queue assignment

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
EDC query management	Detects discrepancies; drafts query text with context; batches for DM review — no direct EDC writes by agents	**T2**	Human DM approval required; controlled query templates; no AI-initiated database changes	Query cycle time; DM review burden; database lock timeline
Deviation and CAPA case management	Standardizes description; proposes severity; routes to CAPA; tracks to closure with recurrence monitoring	**T3**	QA approval required; no auto-closure; severity rubric enforced; recurrence monitoring active	Time to CAPA initiation; classification accuracy; recurrence rate on closed deviations
Submission readiness checker	Pre-validates packages for completeness and cross-reference integrity; routes gaps before deadline	**T3**	No auto-submission; controlled repository only; RA sign-off on all gap resolutions	Right-first-time rate; RFI count; submission cycle time
Labeling change impact mapper	Traces label changes to all downstream materials across markets and channels; creates remediation task list per affected asset	**T2–T3**	Asset inventory required; human approval on assessment; no auto-update of materials; multi-market tracking	Time from label change to material updates; compliance finding rate; materials updated per cycle
Contract compliance monitor	Continuously checks term adherence across active payer agreements; flags anomalies for market access / finance review	**T3 (alerts)**	Threshold tuning; human validation of exceptions; dual control on financial corrections; SOX controls active	Leakage detected; exception resolution time; error rate on rebate calculations
Change control orchestrator	Maps changes to impacted processes, documents, and validation states; routes approvals; tracks implementation	**T3–T4**	SoD; dual approval for critical changes; validated tool catalog; immutable audit trail	On-time CC completion; impact assessment accuracy; validation finding rate post-change
Site risk monitoring and RBM orchestration	Continuously scores active sites on enrollment lag, protocol deviation rate, query age, and document completeness; routes targeted monitoring plan updates and CRA escalation tasks	**T3**	Conservative risk thresholds; CRA management review before any monitoring plan change; no automated site suspension; SLA-governed escalation; validated risk model	On-site monitoring cost reduction; high-risk site identification lead time; protocol deviation detection rate vs. standard monitoring
Safety signal triage and routing	Applies disproportionality analysis (ROR/PRR) against accumulating case data; flags emerging signal candidates; routes to PV physician with structured evidence package for medical review	**T3**	Validated statistical thresholds; PV physician review required before any signal enters formal evaluation; immutable case audit trail; no agent-initiated regulatory notification under any tier	Signal detection lead time vs. aggregate review cycle; false positive rate; on-time PSUR contribution rate

Level 4: The Role-Based Operating Layer

> ***Level 3 gives the enterprise the operational spine. Level 4 is how the enterprise feels unified to every expert who works within it — eliminating system-hopping and embedding institutional knowledge into the daily work surface of every function.***

Pharmaceutical enterprises run on specialized expertise operating across a deeply fragmented application landscape. A CRA navigates CTMS, eTMF, EDC, regulatory portals, and email to complete a single monitoring follow-up. A PV processor moves between the safety database, medical literature tools, prior case records, and reporting templates. A regulatory writer assembles a HA response from RIM, the prior submission archive, the data warehouse, and the scientific record — in separate applications with no shared context. An MSL prepares for a KOL meeting from CRM, approved slide library, published literature, and prior interaction notes — assembled manually before every call. Each context switch taxes expert attention, creates an opportunity for inconsistency, and delays the work that requires actual judgment.

LEVEL 4 The Role-Based Operating Layer

Unified, policy-enforced surface for every function — tasks, evidence, and governed actions in one place, institutional knowledge embedded

GxP Role Surfaces

- CRA copilot: CTMS task queue, site status, protocol deviations, prior findings, eTMF completeness status, and MVR drafting surface in a single interface with policy-aware action buttons and tier-enforced boundaries
- PV case processor workspace: intake queue, narrative drafting surface, follow-up dispatch (T4 after approval), MedDRA coding suggestions (T2), SLA tracker, and prior case search — unified; no cross-system navigation required to complete a case
- Regulatory writing environment: submission history, prior response language (version-pinned), controlled data package, and draft workspace with mandatory citation tracking and claim-checking integrated into every drafting action
- Quality operations surface: deviation queue, SOP reference (version-pinned from Level 2), CAPA tracker, EM trend digest, severity classification support, and change control status — one

Commercial, Medical & Market Access Surfaces

- MSL pre-call intelligence: KOL profile, prior interaction notes, approved scientific content, recent KOL publications, and compliance-checked engagement plan — assembled automatically from governed sources; no manual pre-call aggregation
- Medical information workspace: MI database search, approved label language (version-pinned), published data citations, and draft response surface with promotional compliance checks — medical reviewer approves before any communication
- Commercial territory surface: account status, formulary position, pull-through metrics, approved content library, and compliance-flagged field activities — territory and HCP restriction enforcement active on every recommended action
- Market access dossier view: formulary status by payer, contract performance tracking, HEOR evidence library (version-pinned), and

What Level 4 Makes Possible at Level 5

- Well-defined role-based surfaces with documented COU make the scope of safe agent automation clear: when a workflow is codified at Level 4, the T3/T4 agent actions that can safely operate within it are identifiable, documentable, and validatable
- Onboarding time for new staff decreases significantly: the operating surface contains the institutional knowledge that previously required years of system navigation experience to internalize, compressing time-to-productivity for CRAs, PV processors, quality associates, MSLs, and market access specialists
- Cross-site and cross-function variation narrows as all professionals in a given role operate through the same policy-enforced surface with the same version-pinned knowledge base — producing more consistent outputs and fewer site-to-site compliance variations

governed view for the QA associate ▪ Manufacturing cockpit: batch review package, anomaly flags, specification cross-checks, EM monitoring status, and deviation routing — all within GMP-governed boundaries; human approval gates preserved and architecturally enforced	submission deadline tracker — unified for the access team with contract compliance status visible ▪ MLR coordinator surface: asset pipeline, review status by asset, compliance flag summary, version history, and content expiry tracking — reducing the coordination overhead that currently slows the promotional content cycle	▪ Audit trail quality improves continuously as every action through the operating layer is logged, attributed, and linked to the policy version that governed it — without requiring separate documentation by the professional

Level 5: Bounded Autonomous Agents

> ***Deploying agents without measurement is not autonomy — it is unmonitored automation. Level 5 is inseparable from the AgentOps telemetry that proves performance, justifies tier escalation, and provides the evidence that makes agent deployment defensible at inspection.***

Levels 0 through 4 build the infrastructure that makes autonomous AI action in pharma defensible: governance architecture, evidence memory with ontology management, orchestrated workflows with traceable evidence, and role-based operating surfaces with documented context of use. Level 5 deploys AI agents that execute scoped, governed actions without requiring human initiation of every individual step — while maintaining the audit trail, performance telemetry, and human oversight that the regulatory environment requires. Every tier assignment is backed by performance evidence. Every tier escalation is a change control event.

LEVEL 5 Bounded Autonomous Agents

Scoped agents execute governed actions within tier-defined policy caps — every action audit-ready, every tier assignment evidence-based, AgentOps telemetry continuous

Agent Deployments by Function

- PV intake triage agent: extracts structured fields, classifies by seriousness/expectedness, de-duplicates, routes ICSRs at T3 — narrative draft generated at T2 for processor completion and electronic sign-off
- PV follow-up orchestration: generates follow-up questions from case gaps at T2, dispatches via approved templates after processor approval at T4, tracks responses against GVP reporting clocks
- eTMF monitoring agent: scans continuously against document inventory, drafts corrective task notifications at T3, tracks resolution — no agent-initiated TMF filings under any tier
- Deviation triage agent: standardizes description, proposes severity with rubric reference, drafts CAPA scaffold at T2–T3 — human root cause analysis required before any CAPA closure
- Batch review support agent: reads MES/eBR data, flags anomalies, assembles structured review checklist at T2 — QA reviewer signs before batch disposition; agent does not determine disposition
- MLR pre-screen and routing agent: flags high-risk content,

Control Architecture (Non-Negotiable)

- Every agent action logged with actor identity, tool called, inputs, outputs, rationale, and policy version — replayable transcript meeting ALCOA+; tamper-evident storage; retained per the applicable document retention schedule
- Tool allow-lists enforced by the policy engine at runtime: actions outside the allow-list are blocked by the engine, not flagged after the fact; allow-lists are version-controlled and updated through formal change control
- Human approval checkpoints for every T3+ action: structured approval UI with mandatory reason codes; sampling audit protocol to detect rubber-stamping (random selection plus frequency-triggered selection for high-volume workflows)
- Kill switches predefined and tested quarterly: automatic escalation to human review if intervention rate drops sharply without

The Expert Capacity Multiplier

- When PV intake agents handle structured triage at T3, case processors concentrate on complex medical assessment and narrative quality — the judgment work that differentiates the function and cannot be delegated
- When quality agents triage deviations and prepare CAPA scaffolds at T2–T3, QA staff focus on root cause reasoning and systemic quality improvement rather than documentation production
- When MLR routing agents pre-screen and route assets at T3, MLR reviewers focus attention on content that actually requires expert judgment rather than on administrative package management
- Industry benchmark data consistently finds 40–60% cycle time compression on high-volume structured case work when bounded agents operate under appropriate governance — translating to pipeline capacity growth without

checks reference completeness, routes to correct reviewer team at T3 — human MLR decision authority preserved on all content ▪ Contract compliance monitoring agent: checks term adherence continuously across active payer agreements, flags anomalies for market access/finance at T3 — dual control on all financial corrections; SOX controls active ▪ Supply chain exception agent: classifies exceptions, proposes allocation scenarios at T2–T3 — human approval required before any supply system action regardless of scenario confidence	corresponding quality gain; circuit breakers if tool failure rate exceeds defined threshold ▪ Tier escalation is change control: moving any agent from T2 to T3, or T3 to T4, requires documented performance evidence from the AgentOps framework, updated COU, risk assessment, and QA sign-off	proportional headcount cost

How You Know It Is Working: The AgentOps Framework

AgentOps is not a dashboard add-on — it is the continuous measurement, monitoring, and evidence infrastructure that keeps agents performing within their validated context of use and provides the empirical basis for every governance decision. In a regulated industry, these metrics are the evidence that justifies the enterprise's operating choices to an inspector who may ask about them. They are also the mechanism that prevents the silent performance degradation that turns a successful Level 5 deployment into an inspection finding.

Metric	Why It Matters in Pharma	How to Measure	What Good Looks Like
Intervention rate	High rate = immature autonomy; sudden drop without quality gain = equally dangerous — both require investigation	Approve / override logs per agent per workflow	Pilot phase 30–70% acceptable; alert if drops >20 pts without matching quality evidence; trend downward over time
Critical safety triggers fired	Confirms guardrails are functioning AND reveals where policies are too loose (not enough triggers) or too tight (legitimate actions blocked)	Policy engine logs with severity classification; review both trigger and no-trigger populations	Alert on any 'high severity' trigger; root cause within 48 h; zero patient-safety triggers should be accepted without investigation
Hallucination / grounding failure rate	PV narratives, regulatory responses, and MLR content must be evidence-based — ungrounded output is an inspection risk and a patient safety risk	Eval harness with golden reference sets; citation audit sampling per domain	Near-zero for GxP-touching outputs; block tier escalation if rising; track separately for each functional domain
Rework rate (major revision required)	The leading quality indicator — tells you whether the AI is genuinely helping experts or creating work to undo	Reviewer feedback tags on AI-assisted drafts; mandatory major/minor classification	Downward trend quarter-over-quarter target; plateau = model or prompt refresh needed; alert if rising in any function
Audit reconstruction success rate	If a run cannot be reconstructed from the	Quarterly drill: reconstruct 20+ sampled	>99% reconstructable; any failure triggers root cause and tier review;

Metric	Why It Matters in Pharma	How to Measure	What Good Looks Like
	transcript it cannot be defended at inspection — this is your governance proof metric	agent-assisted decisions from transcript alone; score completeness	include cross-function sample in every drill
Data leakage incidents	Zero tolerance for PII, patient/subject personal data, or trade-secret content in logs, outputs, or external model training feeds	DLP monitoring + incident reports + quarterly log sampling	Target: 0 confirmed incidents; any suspected incident triggers agent pause, forensic review, and regulatory notification assessment
Model drift indicators	Validated performance must remain stable — drift without detection is a silent inspection risk	Periodic re-evaluation against baseline test sets; production metric monitoring on rolling basis	Alert on statistically significant performance drop vs. acceptance baseline; re-validation required before continued operation
Tool failure rate	Drives reliability gaps and incomplete audit trails — a failed tool call is an unrecorded action in the evidence chain	Distributed tracing; tool-call success/failure logging per invocation	Alert if >1–2% sustained; >5% triggers immediate investigation; failure mode analysis per tool required
Cost per completed task	Keeps the program sustainable — token, compute, and integration costs must be governed as the agent estate scales	FinOps tagging by agent, use-case class, and tier	Budget caps defined at program start; alert if spikes >25% vs. baseline without scope change; report monthly to program sponsors

The Reconstruction Audit

The most important periodic check in the pharma AI program: quarterly, select a sample of agent-assisted decisions across functions — PV cases, deviation records, submission artifacts, commercial compliance actions, contract exception resolutions — and reconstruct each one end-to-end from the immutable transcript alone. Target: more than 99% of sampled decisions should be fully reconstructable. Who acted. Which tool. What input. What output. What rationale. Which policy version. Any case that cannot be reconstructed to this standard is an inspection finding waiting to materialize. The discovery of even one non-reconstructable case triggers immediate root cause investigation and temporary tier reduction for the affected agent until the evidence gap is fully resolved.

Level 5 Use Cases by Domain

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
PV intake triage agent	Extracts structured fields; classifies by seriousness/expectedness; de-duplicates; routes to correct queue	**T3**	SLA rules; escalation for high-severity; human medical assessment required; full case audit trail	On-time ICSR reporting rate; triage accuracy; time from receipt to queue assignment
PV follow-up dispatch	Drafts follow-up questions from case gaps; dispatches via approved templates after processor approval	**T4**	Dual approval; approved templates only; GVP clock tracking; T4 requires QA sign-off on tier	Follow-up completeness; response turnaround; GVP SLA compliance

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
eTMF monitoring agent	Continuously scans against expected document inventory; drafts corrective task notifications; tracks resolution	**T3**	No auto-filing; controlled vocabulary for classification; QC sampling on AI classifications	TMF completeness %; gap resolution cycle time; inspection finding rate
Deviation triage agent	Standardizes description; proposes severity with rubric reference; drafts CAPA scaffold	**T2–T3**	QA approval before any issuance; no auto-closure; rubric enforced; recurrence monitoring	Time to CAPA initiation; classification accuracy vs. QA final; recurrence rate
Batch review support	Reads MES/eBR data; flags anomalies and spec deviations; assembles structured review checklist	**T2**	Human QA sign-off before any disposition; agent does not determine disposition; GMP audit trail	Batch review cycle time; anomaly capture rate; release quality metrics
MLR routing agent	Flags high-risk content; checks reference and fair-balance completeness; routes to correct reviewer team	**T3**	No AI clearance of content; human MLR authority preserved; dual approval for final release	MLR cycle time; rejection rate at formal MLR; compliance finding rate on released assets
Contract compliance agent	Checks term adherence continuously across active payer agreements; flags exceptions for resolution	**T3**	Threshold tuning; human validation of exceptions; dual control on financial corrections; SOX controls	Leakage detected; error rate on rebates; exception resolution cycle time
Supply exception agent	Classifies supply exceptions; proposes allocation scenarios; coordinates stakeholders	**T2–T3**	Human approval before any supply action; policy caps on scenario scope; SLA monitoring	Fill rate; expedite cost trend; exception resolution time; service level compliance
Stability trend monitor	Detects OOT trends in LIMS data; proposes investigation scope; escalates to quality team	**T2–T3**	QA oversight on all escalations; human investigation decision required; conservative trend thresholds	OOT detection lead time vs. manual; OOS prevention rate; investigation initiation time
KOL mapping agent	Analyzes publication records, trial participation, and speaking activity to map influence networks	**T2**	Medical director review of all outputs; privacy controls; no commercial use of medical engagement data	Identification accuracy; engagement plan quality scores; MSL adoption
RBM risk scoring agent	Continuously re-scores all active sites across all trials; triggers targeted monitoring plan updates and SDV scope adjustments within validated policy caps; escalates high-risk sites to CRA management	**T3**	Policy caps on SDV scope changes; CRA management approval for all escalations; validated risk model with re-validation on schedule; full audit trail on all scope changes; kill switch active	On-site visit cost reduction; high-risk site detection lead time vs. reactive monitoring; inspection finding rate

Level 6: The Predictive Value Chain

> ***The most expensive events in drug development are the ones that were visible in advance but not acted on: the enrollment shortfall that became a trial delay, the manufacturing drift that became a batch failure, the safety signal that accumulated undetected across two aggregate periods. Level 6 converts operational data accumulated at Levels 3 through 5 into forward-looking signals that allow the enterprise to act while intervention is still low-cost.***

By the time most pharmaceutical operational failures are visible in a status report, the low-cost intervention window has closed. Level 6 addresses this through continuous modeling of operational, scientific, and commercial data. Critically, the accuracy of Level 6 models is directly proportional to the quality of the data foundation built at Levels 2 and 3. Organizations that underinvested in Level 2 ontology management consistently find Level 6 models unreliable — the fix is not a better model, it is returning to complete the ontology work that was deferred.

LEVEL 6 The Predictive Value Chain

Risk, quality, signal, and commercial data modeled continuously — ontology conflicts surface as anomalies — leaders manage proactively

Clinical & Safety Prediction

- Clinical trial risk prediction: enrollment trajectory, site performance, and protocol deviation patterns modeled to flag timeline risk 4–8 weeks earlier than retrospective reporting — enabling proactive site remediation before amendment is required
- PV signal emergence modeling: disproportionality statistics, case volume patterns, and literature signals combined to surface potential safety signals earlier in the aggregate evaluation cycle — earlier benefit-risk assessment with less accumulated patient exposure
- SAE and reporting timeline forecasting: reporting clock risk and aggregate period deadline monitoring with early-warning alerts for functions approaching compliance thresholds
- Site health scoring: composite index updated continuously to identify early indicators of data quality, recruitment, or protocol compliance decline before

Quality, Manufacturing & Regulatory Prediction

- Manufacturing quality drift detection: SPC on LIMS and historian data surfaces OOT trends before they produce OOS results or batch failures requiring regulatory notification — the highest financial-return prediction in the manufacturing function
- Inspection risk scoring: continuous assessment of eTMF completeness, cross-reference integrity, and evidence quality against historical inspection finding patterns — pre-inspection preparation guided by actual risk, not generic checklists
- Regulatory intelligence synthesis: HA feedback patterns, competitor submission outcomes, and label evolution trends analyzed to inform filing strategy and anticipate agency questions before submissions are filed
- Ontology conflict detection: when cross-function predictions produce contradictory outputs for the same underlying entity,

Commercial, Medical & Market Access Prediction

- Demand and supply forecast intelligence: sales trajectory, competitive entry, and formulary change signals modeled to anticipate supply and inventory requirements ahead of market events
- Territory under-performance prediction: leading indicators identified weeks before revenue impact is visible in lagging sales data — enabling market access interventions that preserve revenue the trailing indicator would never recover
- Contract compliance risk: deviation patterns across thousands of commercial agreements surface leakage risk before it accumulates to material financial exposure; anomalies routed to market access and finance for human action
- Access barrier pattern detection: time-to-therapy bottlenecks and patient journey abandonment

formal escalation becomes necessary	the model surfaces this as an anomaly routed to the Level 2 ontology governance function — Level 2 improves from Level 6 telemetry	points identified from hub and patient support data in near-real time to guide access strategy adjustments

Data Foundation Return: Level 6 is where the Level 2 ontology investment pays its largest return. Prediction models operating on clean, consistently defined, version-pinned data are dramatically more accurate than those operating on fragmented inputs. If Level 6 models are underperforming, the root cause is almost always in the data foundation — not in the model architecture.

Level 7: The Adaptive Enterprise — Where the Data Problem Is Fully Resolved

> ***Level 7 is not the destination for AI tools — it is the destination for an enterprise whose data infrastructure, governance architecture, and operational intelligence have matured to the point where the value chain operates as a continuously improving, self-governing system.***

At Level 7, the pharmaceutical enterprise fully resolves the data foundation problem introduced at the beginning of this document. The most significant operational change is the completion of the unified data fabric: a single authoritative identity layer across all systems, a governed enterprise ontology that maps terminology across GxP, commercial, and scientific domains, a self-updating scientific knowledge graph, and a continuous data quality monitoring function that treats ontology drift as an operational risk metric with defined alert thresholds and mandatory remediation procedures.

LEVEL 7 The Adaptive Enterprise

Unified data fabric, governed enterprise ontology fully resolved, self-updating scientific knowledge graph, and continuous AI governance

The Unified Data Fabric

- Single authoritative identity layer: compounds, patients, sites, investigators, suppliers, and HCPs maintained as a governed entity registry across all systems — the same entity is the same record regardless of which system is queried, eliminating the identity fragmentation problem
- Governed enterprise ontology: formally maintained vocabulary layer resolves terminology across clinical (MedDRA), manufacturing (deviation taxonomy), commercial (therapeutic area), and scientific (target biology) domains into a unified, version-controlled reference
- Self-updating scientific knowledge graph: new publications, regulatory guidance, competitive submissions, and real-world evidence integrated continuously into a structured, citable knowledge resource under governance — not a static

Continuous Value Chain Intelligence

- Integrated development orchestration: protocol, site, data, safety, and regulatory artifacts managed as a single governed case from IND through NDA/BLA with complete cross-function visibility and continuous evidence chain maintenance
- Autonomous quality and manufacturing monitoring: batch review, deviation triage, CAPA tracking, and EM trend analysis operating continuously at T4–T5 within validated scope — human oversight reserved for exceptions and critical decisions
- Continuous PV intelligence: post-market safety integrating case data, literature, and real-world evidence into a synthesized, always-current safety picture with proactive signal management across the full product lifecycle
- Commercial and medical intelligence platform: territory performance, competitive positioning, market access landscape,

Self-Improving AI Governance

- Agent performance, policy effectiveness, and model drift monitored continuously against validated baselines — governance adapts based on operational evidence rather than periodic manual review alone
- Tier escalation and de-escalation decisions are evidence-driven in their recommendation, human-governed in their approval — no agent operates at a tier its telemetry history does not support
- Portfolio governance with AI: pipeline decisions informed by continuous evidence of development risk, competitive landscape, and market opportunity — updated quarterly rather than assembled once per year
- Outcome-based commercial structures become credible commitments: the evidence architecture that Level 7 provides supports risk-sharing agreements, outcome-based payer contracts, and performance-based milestone deals that

repository but a living system ▪ Continuous data quality monitoring: ontology drift, identity fragmentation, and provenance gaps treated as operational risk metrics with defined alert thresholds, named owners, and quarterly remediation audits — the data foundation is never assumed to be stable ▪ Full data lineage from source to prediction: every model output traceable to the specific data points, ontology versions, and analysis parameters that generated it — reconstructable for inspection and reproducible for regulatory submission	and promotional compliance managed through continuous monitoring with adaptive strategy recommendations for leadership review	require exactly the evidence quality and operational predictability the Level 7 enterprise can demonstrate

What Level 7 Actually Requires

Level 7 is not a technology milestone — it is an organizational maturity milestone. What distinguishes enterprises that reach it is the audit architecture built at Level 0, the knowledge and ontology governance built at Level 2, the orchestration discipline built at Level 3, the operating surface standardization built at Level 4, the agent performance evidence built at Level 5, and the predictive intelligence built at Level 6. The data foundation problem cannot be solved at Level 7 — it must be addressed beginning at Level 0 and built deliberately through every level that follows. The ladder is real, and it must be climbed in sequence.

Commercial, Medical Affairs, and Market Access: Full Domain Treatment

Commercial, medical affairs, and market access represent the highest-volume, highest-compliance-scrutiny AI deployment environment in the enterprise outside of pharmacovigilance. The scale of interactions — thousands of HCP engagements per week, hundreds of payer accounts, dozens of simultaneous MLR cycles, continuous medical information requests across global markets — combined with the promotional compliance and off-label communication constraints that govern every output, makes these functions both the most compelling AI opportunity and the one that demands the most carefully designed governance.

> ***Commercial AI that ignores compliance is not a time-saver. It is a regulatory enforcement action waiting to happen. Commercial AI designed compliance-first from the governance layer up is among the highest-value deployments in the enterprise.***

Medical Affairs and Medical Science Liaisons

Medical science liaisons represent the most scientifically sophisticated field force in the enterprise. Their value is in the quality of scientific exchange they deliver — calibrated to each KOL's research agenda, informed by the full breadth of published evidence, and scrupulously compliant with off-label communication boundaries. AI's role is to ensure every interaction is prepared with more depth, better evidence, and greater compliance confidence than manual preparation allows.

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Pre-call intelligence assembly	Aggregates KOL publications, prior interactions, CRM history, recent congress presentations, and approved scientific content into a structured pre-call brief	**T2**	Approved content only; interaction data access-controlled; no prohibited information; HCP privacy controls	Preparation time; interaction quality scores; scientific exchange depth
Congress insight synthesis	Summarizes scientific themes, notable data, and competitive developments; produces cited action items for medical director review	**T2**	Citation discipline enforced; medical reviewer sign-off; commercial insights kept strictly separate from medical outputs	Synthesis accuracy; time saved; medical relevance rating
Medical information response drafting	Drafts balanced, cited response to unsolicited MI request using approved MI database and published literature	**T2**	Approved MI database only; off-label escalation enforced; citation required for every clinical claim; medical approval before sending	Response time; off-label escalation rate; response quality scores
KOL identification and mapping	Analyzes publication record, trial participation, speaking activity, and institutional affiliation to identify emerging KOLs	**T1–T2**	Medical director review of all outputs; privacy controls; no commercial use of medical engagement data	Identification accuracy; scientific exchange quality; engagement plan effectiveness
Publication planning support	Tracks publication timelines, abstract submissions, manuscript versions, and author coordination; flags deadline risks	**T2–T3**	Author attribution compliance required; no ghostwriting; ICMJE guidelines enforced; human	Timeline adherence; manuscript completion rate; author coordination overhead

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
			author sign-off on all content	
Field insight synthesis	Synthesizes MSL interaction reports, advisory board inputs, and congress observations into strategic intelligence for clinical and medical leadership	**T2**	Source-specific insight de-identification where required; human leadership review; no commercial use of medical insights	Insight quality; strategic action rate; cycle time from field to synthesis
Medical education content support	Supports independent medical education with literature synthesis and citation mapping — human scientific author retains full authorship	**T1–T2**	Full author control; no undisclosed AI authorship; independence from promotional intent technically enforced	Content accuracy; citation completeness; reviewer feedback scores

Commercial Operations and Sales Force Effectiveness

Commercial operations at enterprise scale — thousands of HCP interactions per week, territory coordination across hundreds of representatives, simultaneous MLR cycles across therapeutic areas and markets — is an environment where AI creates the greatest operational leverage and the greatest compliance risk simultaneously. The governance track established at Level 0 for commercial functions is what makes deploying AI here an advantage rather than an enforcement exposure.

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Governed next-best-action	Suggests compliant engagement actions and approved content; explains rationale; flags restricted territories and HCP status	**T2–T3**	Approved content library only; territory/HCP restriction enforcement; audit trail on all recommendations; no prohibited content suggestions	Commercial productivity; HCP engagement rate; compliance flag rate on suggested content
Territory briefing pack assembly	Compiles account status, formulary position, pull-through metrics, prior call notes, and approved content	**T2**	PII and personal data redaction on patient-level data; compliant content filter; human review before any customer interaction	Preparation time; adoption rate; call quality scores
MLR pre-screen agent	Flags high-risk language, unsupported claims, missing fair balance, and incomplete references before formal MLR	**T2–T3**	Conservative thresholds; no AI clearance of promotional content; human MLR decision authority preserved	Review cycle time reduction; rejection rate at formal MLR; compliance findings on cleared assets
MLR package orchestration	Assembles review-ready package; checks metadata completeness; routes to correct reviewer team; tracks approval status and expiry dates	**T3**	Checklist enforcement; dual approval for final release; no auto-release of any content; audit trail on every routing decision	MLR cycle time; package completeness at submission; late approval rate
Field inquiry triage to medical information	Routes unsolicited HCP inquiries to appropriate pathway based on content classification; drafts intake notes	**T3**	Conservative off-label detection; escalation to MI for ambiguous cases; audit trail; no AI response to off-label without medical review	Routing accuracy; off-label escalation rate; HCP response time

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Label change impact on promo materials	Maps label and safety communication changes to all affected promotional materials; creates remediation task list	**T2–T3**	Asset inventory required; human approval on impact assessment; no auto-update of materials; multi-market coordination tracking	Time from label change to material updates; compliance finding rate; materials updated per cycle
Content reuse and version management	Identifies reusable approved components; tracks expiry dates; flags content needing re-review due to label or safety changes	**T2–T3**	Controlled content library only; human approval for reuse decisions; automatic expiry enforcement	Content reuse rate; expired content usage rate; version consistency across markets
Sales incentive design simulation	Generates scenario comparisons, attainment distributions, and perverse incentive risk analysis for HR and leadership review	**T1–T2**	Human design authority; fairness review required; SOX controls on compensation data	Scenario quality; design cycle time; attainment distribution outcomes post-implementation

Market Access and Payer Operations

Market access bridges clinical evidence to commercial revenue — translating trial outcomes into formulary positions, HEOR evidence into value dossiers, and payer relationships into access agreements. The AI opportunity in market access is significant and underexplored: the function operates at the intersection of high-stakes evidence synthesis, multi-stakeholder relationship management, and financial contract governance where errors have direct revenue and compliance consequences.

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Value dossier and HEOR assembly	Assembles HTA submissions, HEOR evidence packages, and value story frameworks from clinical data, RWE, and economic models	**T2**	Controlled evidence sources; HEOR team review required; no unvalidated economic models; human approval before any submission	Preparation time; HTA acceptance rate; evidence request fulfillment speed
Tender response assembly	Drafts tender responses with evidence synthesis, competitive positioning, and pricing proposal from approved templates and evidence library	**T2**	Template constraints; approved claims only; medical and legal review before submission; audit trail on all submitted content	Tender cycle time; win rate trend; requirement fulfillment rate
Contract compliance monitoring	Checks term adherence against claims, rebate calculations, and performance commitments across active agreements; flags anomalies	**T3 (alerts)**	Threshold tuning; human validation of exceptions; dual control on financial corrections; SOX controls active	Leakage detected; exception resolution time; error rate on rebate calculations
Payer account intelligence	Compiles formulary status, rebate performance, access barrier analysis, and competitive positioning into structured account intelligence	**T1–T2**	Human market access team review; commercial use separated from medical use; privacy controls on payer data	Account preparation time; intelligence quality; team adoption
Time-to-therapy bottleneck analysis	Identifies patient journey abandonment points and access barrier patterns from	**T2**	Patient data de-identification; human access team decision on all	Time-to-therapy improvement; abandonment rate

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
	hub data and prior authorization records		interventions; privacy controls enforced	reduction; access barrier closure
HEOR real-world evidence synthesis	Synthesizes RWE from published studies, registry data, and retrospective analyses to support formulary and label extension strategy	**T1–T2**	Evidence quality assessment required; citation for all claims; HEOR team validation before any submission use	Evidence synthesis time; evidence package strength; RWE utilization in submissions
Outcome-based contract evidence support	Assembles outcome measurement data and performance evidence for outcome-based payer agreements; tracks against contracted milestones	**T2–T3**	SOX controls on financial data; dual control on performance reporting; human sign-off on milestone evidence packages	Milestone reporting accuracy; contract renewal rate; leakage prevention

Patient Support Programs and Hub Operations

Hub operations handle prior authorization management, specialty pharmacy coordination, reimbursement support, and patient assistance enrollment across thousands of cases simultaneously. AI creates the operational capacity to manage this volume with appropriate human oversight and privacy controls at each consequential step.

Use Case	What the Agent Does	Tier	Key Controls	Value Signal
Case routing and eligibility prep	Routes incoming patient cases by payer type, diagnosis, and program eligibility; assembles preliminary eligibility documentation for hub specialist review	**T3**	Consent enforcement; privacy controls; human hub specialist review before patient contact; audit trail on all assignments	Time-to-therapy; routing accuracy; eligibility prep completeness
Prior authorization support	Tracks PA status across cases; drafts appeal letters from clinical evidence and prior successful appeals; flags abandonment risk	**T2–T3**	Human specialist review of all PA materials; privacy controls; no auto-submission to payers without human approval	PA approval rate; appeal success rate; abandonment rate due to PA delay
Specialty pharmacy coordination	Manages prescription transfer, inventory confirmation, and patient notification workflows with specialty pharmacy partners	**T3–T4**	Template-based communications; human approval on exception resolution; privacy controls; SLA monitoring for time-sensitive shipments	Fill rate; exception rate; time from prescription to dispense
Patient assistance program coordination	Tracks financial assistance eligibility, application status, and enrollment; flags cases requiring case manager attention	**T3**	Privacy-preserving design; consent documentation required; human case manager oversight; audit trail	Enrollment rate; application completion rate; case manager capacity improvement
Time-to-therapy bottleneck monitoring	Identifies drop-off points and access barriers in real time; surfaces patterns for access team intervention; proposes outreach prioritization	**T2**	De-identification; human access team decision on all interventions; equity monitoring on prioritization algorithms	Time-to-therapy improvement; abandonment rate; intervention response rate

Function-by-Function Capability Progression

The table below summarizes the full maturity arc for each enterprise function and the primary financial mechanism that drives return across the progression.

Function	Capability Progression (L0 → L7)	Primary Financial Mechanism
Discovery & Preclinical	L0: Governed AI in ELN & lab systems. L1: Literature synthesis, assay drafts, report scaffolds. L2: Versioned scientific record; compound ontology governed. L3: IND package orchestration; CRO workflow cases. L4: Scientist copilot: protocol + evidence unified. L5: Autonomous pipeline execution; lab scheduling agents. L6: Candidate attrition prediction; target opportunity scoring. L7: Self-updating scientific knowledge graph; continuous target intelligence.	Expert capacity multiplier on scarce scientific talent; faster candidate identification reduces time-to-IND; quality of evidence entering development improves compound attrition rates downstream
Clinical Operations	L0: GCP AI boundary defined. L1: MVR drafts, query letters, ICF support. L2: Current protocols always accessible; site ontology matched. L3: Site activation and eTMF orchestrated end-to-end. L4: CRA unified surface. L5: eTMF monitoring agents; data cleaning agents. L6: Enrollment risk predicted 4-8 weeks early. L7: Integrated trial intelligence from IND to CSR.	Cycle time compression on trial milestones; inspection readiness maintained continuously; expert CRA capacity concentrated on site relationship and scientific work
Pharmacovigilance	L0: GVP boundary defined; safety AI governed. L1: Case narrative drafts; MedDRA coding suggestions. L2: Coding ontology versioned; prior cases citable. L3: Intake routing and follow-up orchestrated to GVP SLAs. L4: PV processor unified workspace. L5: Triage agents (T3/T4); follow-up dispatch. L6: Signal emergence modeled continuously. L7: Continuous post-market safety intelligence.	GVP SLA compliance without proportional headcount growth; signal detection quality improves; aggregate reporting cycles shorten; patient exposure to undetected signals reduced
Regulatory Affairs	L0: AI boundary for submissions and labeling defined. L1: Q&A drafts; dossier gap lists. L2: Prior submissions versioned; controlled claims citable. L3: HA correspondence cases; submission dependency tracking. L4: Regulatory writer unified surface. L5: Submission readiness agents; RIM consistency checks. L6: HA intelligence synthesis; filing risk prediction. L7: Continuous regulatory lifecycle management.	Submission cycle time compression; right-first-time rate improvement; HA Q&A response quality improves; inspection confidence in evidence quality increases
Manufacturing & Quality	L0: GMP AI boundary defined (Part 11/Annex 11). L1: Deviation drafts; CAPA scaffolds; SOP pre-checks. L2: SOPs and batch specs versioned; single source of truth. L3: Change control and deviation case management orchestrated. L4: QA unified surface. L5: Batch review agents; EM trend monitoring. L6: Manufacturing drift detected before OOS. L7: Continuous quality governance; self-updating process intelligence.	Batch release cycle time reduction; inspection finding rate reduction; OOS prevention through early drift detection; change control efficiency improvement
Supply Chain	L0: Supply and SOX AI boundary defined. L1: Shortage scenario narratives; S&OP pack assembly. L2: Supplier specs and contracts versioned. L3: Cold chain exception and chargeback case management. L4: Planner unified surface. L5: Supply exception agents within policy caps. L6: Demand and supply risk predicted ahead of market events. L7: End-to-end supply network intelligence.	Supply reliability improvement; expedite cost reduction; leakage detection and prevention; S&OP decision quality and cycle time

Function	Capability Progression (L0 → L7)	Primary Financial Mechanism
Commercial Operations	L0: Promotional compliance AI defined; OPDP boundary enforced. L1: MLR pre-screen; territory briefing packs; call prep. L2: Approved claims library; label language always current. L3: Labeling change impact on promo assets orchestrated. L4: Sales surface with compliance-enforced content. L5: MLR routing agents; field inquiry triage. L6: Territory under-performance predicted early. L7: Continuous commercial intelligence and compliance.	Commercial team productivity; MLR cycle time reduction; promotional compliance improvement; territory under-performance detected weeks earlier
Medical Affairs	L0: Medical AI boundary; off-label controls defined. L1: Congress synthesis; MI response drafts; MSL pre-call briefs. L2: Approved scientific content versioned; publications citable. L3: Insight collection workflow; publication planning cases. L4: MSL unified surface. L5: Insight synthesis agents; KOL mapping. L6: KOL influence trends; evidence gap identification. L7: Continuous scientific landscape intelligence.	MSL scientific exchange quality; MI response speed and accuracy; publication compliance and timeliness; insights translated to clinical and commercial strategy
Market Access	L0: Market access AI boundary defined. L1: Value dossier drafts; tender response assembly. L2: HEOR evidence library versioned; formulary data citable. L3: Payer account cases; contract compliance tracking. L4: Access team unified surface. L5: Contract compliance monitoring agents; access barrier detection. L6: Formulary shift and access barrier patterns predicted. L7: Outcome-based contract evidence platform; continuous payer intelligence.	Formulary submission quality and cycle time; contract leakage detection and prevention; time-to-therapy improvement; HTA acceptance rate

The Platform Landscape: Where AI Agents Plug In

The pharmaceutical enterprise's AI architecture is built on the integration of a diverse, validated application landscape. Each platform category has its own governance boundary, integration requirements, and validation posture. The Level 0 governance architecture must define the AI boundary for each category; the Level 2 knowledge layer must ensure version-pinned retrieval from each validated repository; and the Level 3 orchestration backbone must create defensible audit trails across every cross-platform action.

Platform Category	AI Agent Insertion Points	Representative Systems	Governance Priority
PV / Safety Platforms	Intake triage; duplicate detection; follow-up orchestration; narrative drafting; literature monitoring	Oracle Argus/Safety One, ArisGlobal LifeSphere Safety, Veeva Safety	Patient safety critical: strictest validation; QC sampling on all outputs; immutable case audit trail; GVP SLA enforcement; trial subject data and personal data controls active throughout
Clinical Platforms (CTMS / EDC / eTMF)	Site activation orchestration; eTMF QC and monitoring; query drafting; supply scenario planning	Medidata Rave/CTMS, Oracle Clinical One, Veeva Vault Clinical	Validated environments; multi-partner integration; no direct EDC writes by agents; GCP audit trail on every action
Regulatory / RIM / Labeling	Consistency checks; HA correspondence triage; submission readiness; labeling impact assessment	Veeva RIM, Ennov, Extedo	Controlled vocabulary; single source of truth; traceability to submission of record; controlled document access only
Quality Management (QMS / CAPA)	Deviation/CAPA drafting; change control orchestration; audit prep; SOP compliance checks; batch review support	MasterControl, Sparta TrackWise Digital, Veeva Quality	GMP scope; SoD enforcement; immutable audit; agent suggestions reviewable and traceable; e-signature preserved for human decisions
Manufacturing / MES / Historian	Batch review support; deviation evidence assembly; stability monitoring; scheduling support	Werum PAS-X, Siemens Opcenter, OSIsoft/AVEVA PI	High validation burden; agents support review, not execute critical steps without approval gates; plant network constraints must be respected
Commercial CRM & Engagement	Pre-call context agents; governed next-best-action; compliant content retrieval; call note automation	Salesforce Life Sciences Cloud, Veeva CRM, IQVIA OCE	Promotional compliance enforcement; territory/role enforcement; consent and interaction audit trail; no off-label content generation
MLR / Promotional Content	Pre-screen risk flagging; package assembly and routing; content reuse identification; label change impact	Veeva Vault PromoMats, content DAMs	Conservative thresholds; no AI clearance of content; human MLR authority preserved; full audit trail on all review decisions
Supply Chain Planning	Exception agents; scenario generation; plan narrative; shortage management proposals; S&OP prep	SAP IBP, Kinaxis RapidResponse, Blue Yonder	Human approval on allocation/supply decisions; SOX controls on financial impacts; dual control on shortage management actions
R&D Informatics (ELN / LIMS)	Experiment copilots; protocol drafting; sample tracking; evidence assembly; QC anomaly detection	Benchling, Dotmatics, LabVantage	GLP controls where applicable; instrument integration validation; controlled template libraries; scientific sign-off required

Platform Category	AI Agent Insertion Points	Representative Systems	Governance Priority
Integration / iPaaS	Tool catalog backbone; event-driven triggers; API governance; connector versioning; rate limit management	MuleSoft, Boomi, Informatica	Integration changes must be governed; connector versioning essential for agent determinism and audit trail integrity
Orchestration / Case Management	HITL checkpoints; approval routing; SLA management; multi-agent coordination; exception escalation	ServiceNow Now Assist, Appian, Pega	Control plane: approvals, SoD, immutable audit, kill switch for all agent actions
Data / AI Foundation	RAG pipelines; eval harnesses; model monitoring; private model hosting; feature stores; agent runtimes	AWS Bedrock, Azure OpenAI, Databricks, Snowflake Cortex	Validated environments; private endpoints; data residency enforcement; change control on all model and prompt updates; drift monitoring

How to Sequence the Investment

The maturity model is designed for deliberate progression under the constraints of a regulated industry. The following principles govern sequencing decisions across all eight levels — reflecting both the GxP requirements that make some sequences non-negotiable and the operational realities that determine where early investment generates the most durable return.

Governance Before Features

The single most common failure mode in pharma AI investment is deploying capability before establishing governance. Level 0 must be in place before any AI system touches a GxP or commercial compliance workflow. Every Level 5 agent action, every Level 6 model output, and every Level 7 orchestration artifact is only defensible if the governance foundation is solid. Organizations that deploy first and govern later consistently encounter the same sequence: promising pilot, escalating exception, management decision to pause, costly remediation before the program can continue.

Deploy Level 1 Broadly and Measure Rigorously

AI assistance in expert workflows delivers visible, credible value quickly across the enterprise. More critically, it generates the quality benchmarks — rework rates, reviewer override rates, source grounding metrics, intervention rates — that provide the empirical foundation for every tier escalation decision that follows. Organizations that deploy without rigorous measurement build on an anecdotal foundation. Those that measure from day one build on evidence.

Build Level 2 Before Promising Level 5 or 6

The most expensive failure in pharma AI is deploying autonomous agents or predictive models on fragmented knowledge and discovering that their outputs are unreliable or indefensible under examination. Level 2 knowledge governance — document version control, controlled source repositories, citation requirements, provenance tracking, and first-generation ontology management — is the prerequisite for every capability above it. Organizations that underinvest here consistently encounter the same failure: an agent that performed well in testing degrades in production because its knowledge base contained stale, conflicting, or uncited content that was never discovered during the pilot.

Run Commercial Governance in Parallel, Not After

Commercial, medical affairs, and market access functions require their own governance track running in parallel to the GxP track from Level 0, not subordinate to it or sequenced after it. Organizations that wait to govern commercial AI until the GxP program is complete lose a year of commercial value and deploy with higher risk. Those that design commercial compliance governance as a parallel track from the beginning deploy commercial AI at scale without the enforcement actions that characterize poorly governed commercial deployments.

Earn Autonomy Through Demonstrated Performance

Every tier escalation — from T2 draft to T3 bounded execution, from T3 to T4 material workflow steps — requires documented performance metrics from the current tier, a risk assessment, updated COU documentation, and QA approval. An inspector who sees a well-documented escalation record is satisfied. One who sees an undocumented autonomy increase finds an undocumented computerized system change. The evidence is the governance.

The One Governing Rule

In a regulated industry, the value of every AI capability is directly proportional to the defensibility of the evidence it produces. Governance architecture, knowledge quality, workflow orchestration, and operational infrastructure are not prerequisites to be minimized — they are the conditions under which every AI investment produces returns that compound rather than erode. Organizations that invest in these foundations rigorously consistently outperform those that attempt to shortcut them — not in the first quarter, but over the five-year horizons that determine competitive position in drug development.

The Financial Architecture of AI Value

A common error in pharmaceutical AI investment planning is treating financial value as a single pool of efficiency savings. The actual financial architecture of an AI maturity program in pharma is structured across seven distinct mechanisms, each operating at different levels and each generating value through a fundamentally different path.

Levels 0–1: Reduced Cost of Expert Time on Assembly Work

Documentation efficiency, narrative drafting, and evidence assembly compression translate directly to cost reduction and expert capacity reallocation across every function simultaneously. A 30–50% reduction in the time required to produce a first-quality draft, applied across CRAs, PV processors, regulatory writers, quality associates, MSLs, and commercial operations teams, represents a meaningful reduction in the per-unit cost of every output these functions produce. These returns are measurable within one to two quarters of deployment and fund the infrastructure investment required at Levels 2 and 3.

Level 2: Reduced Rework and Inspection Remediation Cost

Authoritative, versioned knowledge reduces citations of outdated documents, re-decisions based on superseded policies, and the retrospective documentation effort that consumes significant pre-inspection resources. These are risk reductions with direct financial value in an industry where a Warning Letter, Complete Response Letter, or consent decree can cost multiples of the annual AI program investment. The ROI of Level 2 is most visible in what does not happen.

Level 3: Throughput Value and Evidence Integrity ROI

Closed-loop orchestration converts the time between decision and cross-system execution into speed across every consequential milestone: site activation, database lock, batch release, submission filing, labeling update, contract renewal. In a therapeutic area where the daily cost of delay on a late-stage asset runs into seven figures, orchestration that reduces cycle time by even a few days per milestone represents compelling financial return.

Levels 4–5: The Expert Capacity Multiplier

When CRAs, PV processors, regulatory writers, quality associates, MSLs, and market access professionals operate through governed role surfaces with AI handling the assembly and routing work, the effective capacity of the team increases without proportional headcount growth. A PV function that processes the same case volume with 30% fewer processors, or a regulatory function that produces submission packages 40% faster with the same headcount, is not merely more efficient — it is structurally more competitive in an industry where expert talent is the binding constraint on pipeline throughput.

Level 6: Prevention Premium on Development-Stage Risk

The prevention premium in pharmaceutical development is extraordinary. A trial enrollment shortfall identified four weeks earlier changes timeline by months — avoiding CRO renegotiations, site expansion costs, and milestone delay payments. A manufacturing OOT detected before it produces OOS avoids a batch failure, regulatory notification, potential supply disruption, and downstream commercial impact. A safety signal identified one aggregate period earlier enables benefit-risk management with one less period of accumulating patient exposure. These prevention values are pipeline economics-level returns on the predictive investment.

Level 7: Structural Competitive Advantage

The enterprise that operates an auditable, self-improving AI value chain achieves a structural advantage in development cycle time, inspection confidence, commercial execution speed, and partnership attractiveness that late movers cannot shortcut. The compound effect of operating faster, with stronger evidence, under better governance, and with greater regulatory confidence, defines competitive standing over five to ten year horizons. It also unlocks commercial structures — outcome-based contracts, performance-based milestone deals, risk-sharing payer agreements — that require exactly the evidence quality and operational predictability that a Level 7 enterprise can credibly demonstrate.

The Enterprise Imperative

Biotech and pharma enterprises that deploy AI as a collection of departmental tools — without shared governance, evidence architecture, or value-chain integration — will achieve incremental productivity gains while accumulating inspection exposure, integration debt, and the kind of organizational AI fatigue that follows a well-publicized pilot failure. Those that build AI as a governed operating system for the value chain — evidence-first at Level 2, orchestrated at Level 3, unified at Level 4, autonomous at Level 5, and predictive at Level 6 — will achieve something categorically different: an enterprise that discovers faster, develops more reliably, operates under a continuously maintained evidence record, and commercializes with the compliance confidence that meets the highest regulatory bar in the world.

The gap between these two outcomes will compound over five to ten years. The governance architecture, knowledge infrastructure, operational backbone, and organizational confidence built through disciplined maturity progression are capabilities that cannot be shortcut by a late-moving competitor. The data ontology work that begins at Level 2 and resolves at Level 7 cannot be skipped — it must be built, and building it correctly the first time is far less expensive than reconstructing it under regulatory pressure. The maturity ladder describes how this advantage is built, and the pharmaceutical industry's regulatory framework ensures that those who build it correctly cannot be easily overtaken by those who try to skip the foundations.

> ***Every function more capable. Every handoff more traceable. Every decision more defensible. Every inspection more confident. Every commercial action more compliant. This is what AI does for a pharmaceutical enterprise — when it is built as a governed value chain rather than a collection of departmental experiments.***

Published by Verit Global Labs · February 2026 · Framework applies across biotech, pharma, CRO, and CDMO

Not adapted to any specific organization · Contact hero@veritglobal.com for a calibrated engagement