LIFE SCIENCES EDITION

AI that *Earns* its Place

How to make AI deliver results in pharma, biotech and medtech

A field guide to the five stages of AI maturity

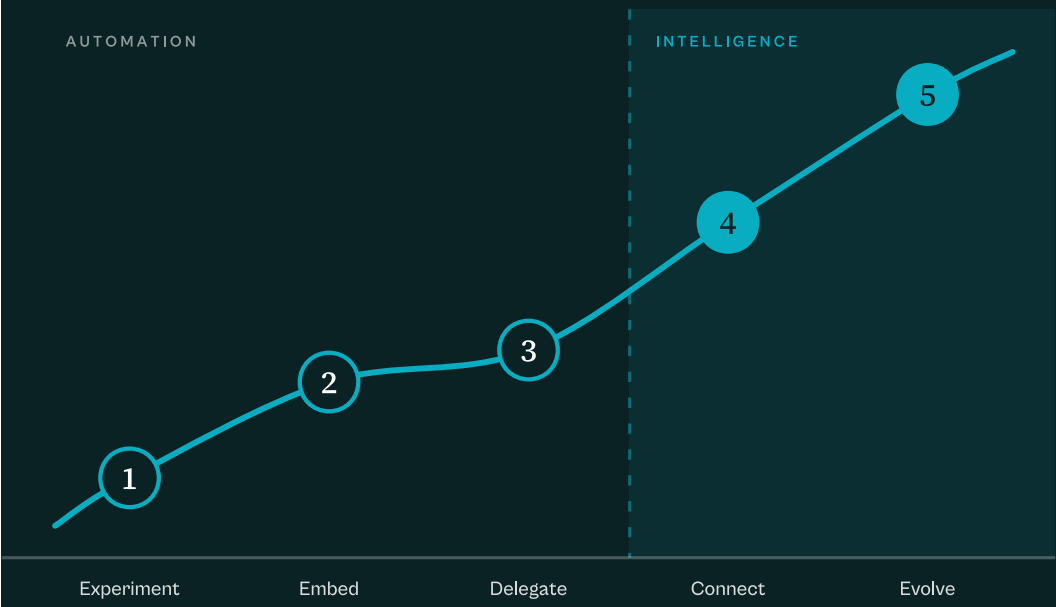

accellor

INSIDE THIS GUIDE

Contents

Most life sciences leaders have seen AI pilots. Fewer can say where their company actually stands, what an AI system should be allowed to do on its own, and what to do next. This guide gives you a plain way to answer all three.

About this guide 3

PART ONE · THE MODEL

Why adoption stalls 7

The five stages of AI maturity 9

One rule: autonomy is earned, and capped by consequence 16

What is actually agentic? 17

Data readiness and the four foundations 18

The software-lite enterprise 21

PART TWO · LIFE SCIENCES

Four examples, end to end 23

PART THREE · NEXT STEPS

Self-check 33

Your first 90 days 35

Sources and notes 36

HOW TO USE IT

Read the opening section and Part One in one sitting. They give you the language.

Then take the self-check (page 33) and the 90-day plan (page 35) to your leadership team for their input. The self-check takes about 20 minutes.

Three questions run through the whole guide: **Where are we? What is this thing? What may it do?**

About this guide

If you read nothing else, read these opening pages and the self-check and 90-day plan at the back.

Companies have run AI pilots for years, and most have not changed how the business operates. The gap is not the technology. Leaders lack a shared way to say where they stand, what an AI system should be allowed to do, and what to implement next. Gartner expects over 40% of agentic AI projects to be canceled by the end of 2027, citing cost, unclear value and weak risk controls.[1] This guide offers a way to avoid being one of them.

The AI Maturity Curve

Five stages, from Experiment to Evolve, resting on four foundations

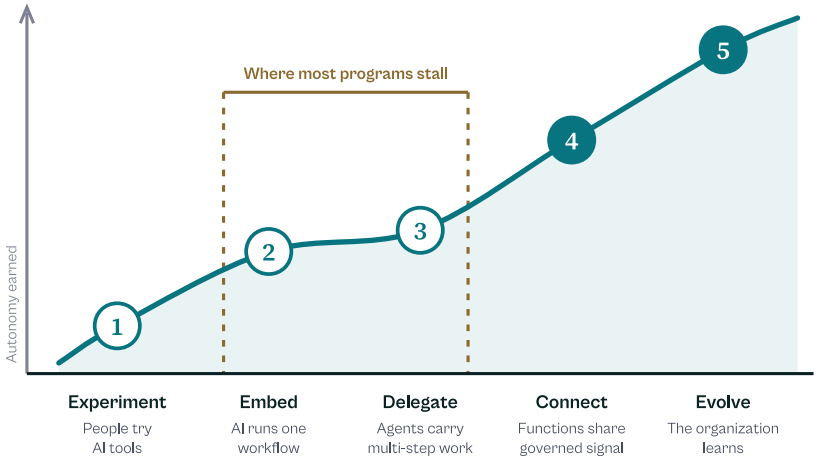

FOUR FOUNDATIONS UNDER EVERY STAGE

Data & knowledge	Governance & assurance	Platform & reuse	People & roles

Most companies are stuck in Stages 2 and 3. This guide gives you a practical framework for getting past the stall, and for capturing the value that comes when truly agentic AI is built into a connected, evolving ecosystem.

The secret to unlocking AI

Building the right tools and choosing the right models is not enough. Task agents and point SaaS solutions will only get you to Stages 2 and 3. The real unlock comes when you design AI into your existing workflows and processes.

AI high performers redesign the work, not just use tools[2]

Share that have fundamentally redesigned workflows

Where AI value actually comes from[3]

Share of the effort behind AI value

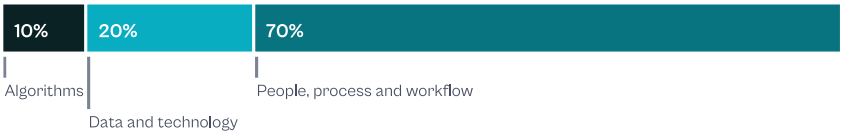

That means AI needs to be:

Customized	to your existing people and processes
Layered	by responsibility: some agents do, some orchestrate and some think
Integrated	into your core systems of record, such as your CRM, ERP and data platforms
Connected	to each other, without the friction of fragmented SaaS point solutions

It frees people to do what they do best: create and decide. As the administrative and execution burden falls, the human workflow changes too. AI value is unlocked only when the AI system and the human system are connected and evolve together.

Stop “agent-washing” your technology

Not every “agent” is agentic. As the technology industry rebrands itself around AI, almost everything is now called an agent. This guide explains the different kinds of AI, from fixed rules to truly agentic systems, and how they must work together (page 17).

AUTONOMY IS EARNED, AND CAPPED BY CONSEQUENCE

A system gets more freedom as it proves reliable, but never more than the cost of a mistake allows. Internal, low-risk work can run with wide limits. Patient safety, product quality and regulatory submissions keep the final decision with a named person.

Five things to remember

1. **Adoption is not a straight line, and it differs by function.** Most programs stall in Stages 2 and 3. Place each function at its own point on the curve.
2. **One governance policy for everything fails.** Match controls to consequence.
3. **Ask what an “agent” really is.** The kind of AI decides how much evidence you need.
4. **People’s work changes, not just their number.** Operational work moves to the system; judgment work grows, and in regulated work it stays with a named person.
5. **Foundations beat pilots.** Shared data, governance, platform and trained people let each new use case ship faster than the last.

WHAT TO DO NEXT

1. Turn to the self-check on page 33 and spend 20 minutes with your leadership team placing each function and checking your foundations.
2. Pick one Stage 2 workflow that carries real expert judgment, and name its owner.
3. Fix the weakest foundation for that workflow, then decide in 90 days what to implement next.

PART ONE

1

The Model

Why do so many AI proofs of concept fail? The AI maturity curve creates a framework for debugging where AI delivers value and where it fails.

Pilots aren't progress. Not every "agent" is agentic. AI autonomy is earned by proof and capped by consequence. A business run by a fragmented ecosystem of a hundred different SaaS solutions is over.

Why adoption stalls

Companies have invested in AI for years, and most of that activity has not changed how the business runs. Three patterns explain the gap. None of them is a technology problem.

Activity without direction

In many large companies, AI pilots multiply faster than anyone can govern them. Each is sponsored locally, built on its own tooling and proven in a demo. Because the pilots share no data layer, no governance model and no common way to evaluate results, the twentieth pilot costs about as much to start as the first. Value shows up as anecdotes rather than outcomes, so leadership cannot tell which pilots deserve to scale. Visible activity is not progress.

One policy for everything

Companies tend to govern AI in one of two ways. Tight, uniform controls block safe, low-risk uses and push people toward unapproved tools. Loose, uniform controls expose quality and patient-safety work to risks nobody has sized. In May 2026, Gartner argued that applying uniform governance across AI agents will lead to enterprise failure. **Controls should rise with autonomy and consequence.** A tool that summarizes meeting notes and a tool that drafts a complaint assessment should not face the same review.

Everything is called an agent

Vendors now use the word “agent” for chatbots, rules engines and genuinely autonomous systems alike. Gartner calls this “agent washing”: the rebranding of assistants, robotic process automation (RPA) and chatbots. In June 2025 it estimated that only about 130 of the thousands of agentic AI vendors are real, and predicted that over 40% of agentic AI projects will be canceled by the end of 2027 because of rising costs, unclear business value or weak risk controls.[1] Buyers who cannot tell what they are buying cannot price the risk.

Why life sciences is different

Generic AI maturity models assume that more autonomy is simply better. Regulated work breaks that assumption, for five reasons.

Constraint	What it means for AI
Validation	Computerized systems used in regulated work must be fit for their intended use, with documented validation and change control (for example, 21 CFR Part 11 and EU Annex 11). A model or prompt change can trigger revalidation.
Accountability	Quality, safety and regulatory decisions are made and signed by named people. An AI system can prepare a decision; it cannot be the accountable party.
Evidence	Decisions must be explainable and auditable, sometimes years later and to a regulator who was not in the room.
Controlled data	SOPs, batch records, case files and complaint records live in validated systems with strict versioning. Agents must respect those boundaries, not work around them.
Global rules	FDA, EMA and other authorities differ in detail. A governance model must flex by region and product.

WHAT LEADERS NEED

- A way to say where the company stands, in terms the board can follow.
- A way to decide what an AI system may do, matched to what is at stake.
- A way to tell real capability from relabeled software.

That is what the rest of this guide provides. It is built for life sciences. A technology can safely draft meeting summaries but should never release a batch or file a submission on its own.

The five stages of AI maturity

Most companies are adopting AI the way they adopted software as a service (SaaS): one tool, one team, one task at a time. Each function gets its own AI, but a finding in one function still reaches the next by email and meetings. This is progress, but it is not transformation.

Accellor describes the journey in five stages. Each stage is defined by a change in who does the work and who is accountable for it, not by how advanced the technology is. In the first three stages, AI does more of the work: this is **automation**. In the last two, functions start sharing what AI learns, and the enterprise learns as a whole: this is **intelligence**.

A company is rarely at one stage everywhere. It may be at Stage 3 in commercial operations and Stage 1 in quality, so place each function separately. The self-check on page 35 shows how.

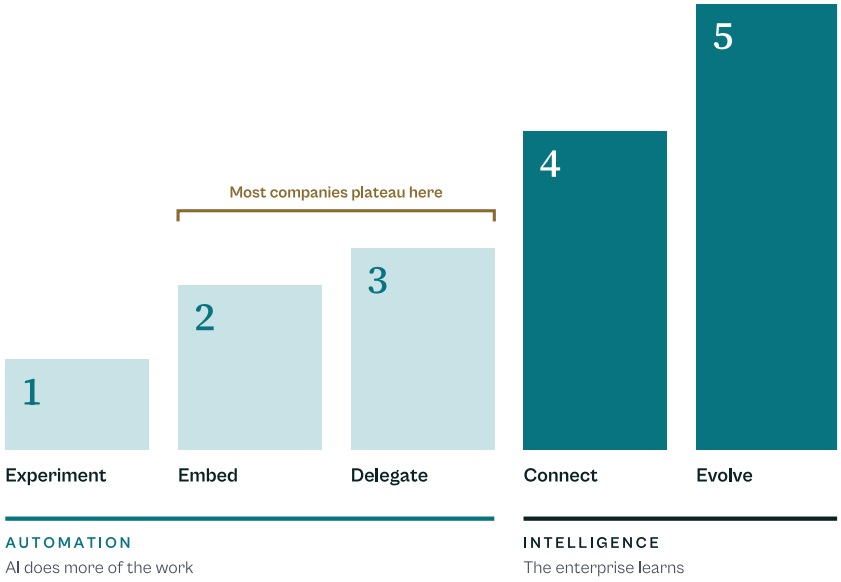

The stages at a glance

Stage	What changes	People's role	Examples (illustrative)
AUTOMATION · AI DOES MORE OF THE WORK			
1 · Experiment	People try AI tools on their own	Early adopter	Drafting emails and notes
2 · Embed	AI runs one workflow; people review	Reviewer	Case coding · Pre-call plans · Defect reports
Where most companies plateau: between Stages 2 and 3			
3 · Delegate	Agents carry multi-step work; experts supervise	Supervisor	Case assembly · Investigation packs · Service visit booking
INTELLIGENCE · THE ENTERPRISE LEARNS			
4 · Connect	Functions share governed signal continuously	Signal steward	Field adverse event to safety · Field failure to plant · Complaint to CAPA
5 · Evolve	The organization learns continuously	Process designer	Case intake retuned from corrections · Call plans redesigned from outcomes · Repairs updated from the field

The pages that follow take each stage in turn: what it looks like in practice, how to tell you are there, what to measure and what to watch for.

STAGE 1 OF 5

1 Experiment

People try AI tools on their own.

Individuals use chat assistants for drafting, summarizing and searching. A few teams run pilots. Nobody owns the picture across the company, and rules for approved tools and data are unclear or missing. The risk is not that people are experimenting; it is that nobody can see what data is going where.

IN PHARMA

Medical writers draft first versions of documents. Clinical operations analysts summarize site feedback. Medical information teams search literature faster.

IN MEDTECH

Marketing teams draft materials for physicians. Engineers query standards and prior designs. Regulatory staff summarize guidance changes.

WHAT TO MEASURE

An inventory of tools in use, by whom and on what data; the share of staff with basic training.

WHAT TO WATCH FOR

Confidential or patient-related data entering unapproved tools. Blanket bans that push use underground.

SIGNS YOU ARE HERE

- Many pilots, no single owner
- Value shows up as anecdotes, not measures
- Legal, IT and Quality disagree on what is allowed

> **To reach the next stage:** Choose one workflow where real expert judgment exists, and set ground rules: approved tools, data rules, and who reviews what.

STAGE 2 OF 5

2 Embed

AI runs inside one defined workflow, and people review the output.

The company picks a high-volume workflow and puts AI inside it. People stop doing the most mechanical work and focus on exceptions and quality. This is the first stage that changes how a team operates, and the first that requires a validation and review approach.

IN PHARMA

Drafting adverse-event case narratives for a safety scientist to review. Drafting responses to medical information requests. Routing and summarizing incoming documents.

IN MEDTECH

Triaging incoming complaints. Drafting investigation summaries for a quality engineer to review. Checking labeling content against source documents.

WHAT TO MEASURE

Cycle time, rework rate, reviewer override rate and error severity, each against the manual baseline.

WHAT TO WATCH FOR

Reviewers who rubber-stamp output. Savings counted before quality is shown. A validation plan that starts only after the build.

SIGNS YOU ARE HERE

- The workflow has a named owner and measured cycle time and quality
- Where required, the workflow is validated
- People review every output, or a defined sample

To reach the next stage: Pick the workflow carefully. One with real judgment is a beachhead; a purely mechanical one is a one-off saving.

STAGE 3 OF 5

3 Delegate

Agents carry multi-step work, and experts supervise.

An agent now gathers inputs, takes several steps and prepares the result. The expert stops assembling and starts supervising, deciding and signing off. This is where roles change, and where the consequence tier of the work starts to determine how far autonomy can go.

IN PHARMA

An agent assembles a regulatory query response package from source documents for a regulatory lead to approve. An agent prepares a pre-call plan for a field representative.

IN MEDTECH

An agent compiles a complaint file, checks it against reporting criteria and drafts the assessment for a person to decide. An agent prepares a supplier-quality review pack.

WHAT TO MEASURE

Share of tasks completed without escalation, exception rate, time to decision and audit-trail completeness.

WHAT TO WATCH FOR

Permissions that creep beyond the approved scope. Supervisors without the time or training to challenge the agent.

SIGNS YOU ARE HERE

- Agents have defined permissions and an audit trail
- Supervisors are trained and accountable by name
- Exception and error rates are tracked and reviewed

> **To reach the next stage:** Build shared data access and shared governance, so each new agent does not need its own.

STAGE 4 OF 5

4 Connect

Agents share governed signal across functions.

Functions begin to inform each other continuously instead of through periodic reports. Each function still owns its decisions. What changes is what each one knows when it decides, and how fast.

IN PHARMA

A safety signal reaches medical, regulatory and commercial teams with an audit trail. A supply disruption informs clinical trial supply planning.

IN MEDTECH

Field complaint trends reach design, quality, service and sales. A manufacturing deviation informs supplier management and release planning.

WHAT TO MEASURE

Time from signal to decision-maker, share of signals with a named owner, and decisions changed by new inputs.

WHAT TO WATCH FOR

Signals that no function owns. One function's data errors spreading to others. Unclear decision rights.

SIGNS YOU ARE HERE

- Someone owns the signal each function receives
- Signal flows are governed, logged and reviewable
- Decision rights have not moved; the inputs have improved

> **To reach the next stage:** Treat signal governance as its own discipline: decide who may publish, consume and correct each signal.

5 Evolve

The organization learns continuously, and work is redesigned around what agents do well.

Outcomes feed back into processes, models and decisions on a governed cadence, and work is redesigned for human-agent teams. Few companies are here. Treat it as a direction, not a near-term target.

IN PHARMA

Safety and real-world signals feed back into protocol design and risk management on a governed, reviewable cycle.

IN MEDTECH

Field performance and complaint data feed back into design controls, supplier requirements and training on a governed cycle.

WHAT TO MEASURE

Share of core processes redesigned around human-agent teams; time from a lesson learned to a process change.

WHAT TO WATCH FOR

Treating Stage 5 as a program rather than a result. Reorganizing before the foundations can carry the load.

SIGNS YOU ARE HERE

- Process owners propose redesigns from agent and outcome data
- Governance adapts faster than annual cycles
- Roles and incentives were rewritten, not just tools

One rule: autonomy is earned, and capped by consequence

A system gets more freedom as it proves reliable, but never more than the cost of a mistake allows. Every task falls into one of three consequence tiers, and each tier sets a ceiling on what AI may do there.

Autonomy ceilings by consequence tier

How far AI may go in each tier, from suggesting to acting inside limits

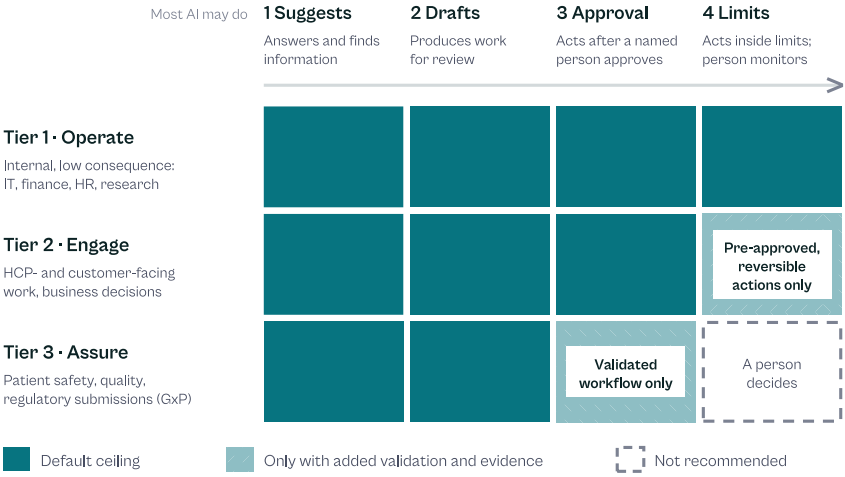

Tier	What it covers	Default ceiling
Tier 1 · Operate	Internal, low-consequence work: IT, finance, HR, research	AI may act inside set limits while a person monitors
Tier 2 · Engage	HCP- and customer-facing work and business decisions	AI acts after a named person approves; acting inside limits only for pre-approved, reversible actions
Tier 3 · Assure	Patient safety, product quality and regulatory submissions (GxP)	AI drafts; acting with approval only in a validated workflow. The decision stays with a named person

Stage and ceiling are different things. Moving up the stages changes who does the work and what each function knows; it does not lift the cap on a Tier 3 decision.

What is actually “agentic”?

Leaders keep asking the same question about every solution: **is this assistance, automation, or truly agentic?**

Four kinds of AI solutions

Kind	How to recognize it	Example
Infrastructure	Fixed rules and code. If the same input always gives the same output and no model decides anything, it is not AI, however it is marketed.	Batch-record routing; validated calculations
Assistant	The model decides content, such as a summary, a score or an answer. A person or calling system decides what happens next.	Literature summary; draft medical-information reply
Bounded task agent	The model owns one narrow goal and chooses its own steps, tool calls and when it is done, inside a defined task.	Assembles a complaint file for one product
Truly agentic	The model plans, selects tools or other agents, recovers from failure and decides when the work is complete. Approval gates bound its autonomy but do not remove it.	Cross-function signal monitoring and response

The further down the table, the more evidence you need before raising authority in Tier 2 and Tier 3 work.

Three questions for any implementation partner creating agents:

1. **Who decides the next step:** your configured rules, a person, or the model?
2. **What happens when it fails partway through a task?** Does it recover, stop or escalate, and who sees the failure?
3. **What can it do without a person approving it,** and how is that limit enforced and logged?

A vendor that cannot answer these precisely is likely selling infrastructure or an assistant under an agent label. That may still be the right product, but it should be priced and governed as what it is.

Data readiness

An agent decides and acts on what it can see. If the data is wrong or disconnected, it acts quickly on a wrong picture of the business. Many deployments plateau at Stage 3 because each workflow's AI works around its own data gaps, and none of them fixes the shared problem.

In a 2026 survey, 34% of life sciences and health care companies rated their data management as highly prepared for AI, a lower share than the year before.[4] In Salesforce's 2026 survey, 96% of IT leaders agreed that the success of AI agents depends on data integration across systems.[5]

What an agent stands on

From the systems of record up to the agents that act on them

Each layer depends on the one beneath it. Tags show which of the five checks below tests it.

Five checks: is this workflow's data ready for an agent?

1	Each customer, product lot and inventory location has one authoritative record.
2	Agents reach that record through governed application programming interfaces (APIs).
3	Permissions define what each agent may read and change.
4	Each data domain has a named owner who corrects errors.
5	Systems publish events, such as a shipment, a scan or a complaint, as they happen.

Data has been central to Accellor's work since before the current wave of AI, from integrating systems and building data foundations to governing master data across CRM, ERP and analytics environments.

Four foundations

Every stage rests on the same four foundations. **The weakest one sets your pace.** Stalled programs are usually held back by one foundation, most often data or governance, rather than by the model. Foundations are also cheaper to build before agents multiply than after.

Foundation	What good looks like	A sensible first move
Data and knowledge	Agents reach trusted, current, labeled documents and records, through access that respects each user's permissions. Controlled documents carry their version and status, and every answer can be traced to its source.	Pick the 20 documents and records your first workflow depends on, and make them reliably retrievable, labeled and versioned
Governance and assurance	Use cases are tiered by consequence. Validation is risk-based and fit for intended use. Changes to models, prompts and data go through change control. Logging, monitoring and an incident process exist before autonomy rises.	Write a one-page tiering rule and apply it to your ten most active AI workflows
Platform and reuse	Teams build on shared components: connectors, orchestration, evaluation tests and logging. The company can change models or vendors without rebuilding workflows.	Choose one approved route for tools and models, and one shared evaluation method
People and roles	Every workflow has a named business owner and a named quality owner. Reviewers and supervisors are trained before deployment. Metrics reward decision quality, not throughput.	Name owners for the beachhead workflow and define the reviewer role before the build starts

Why this order matters

A shared data layer, a tiered governance model and a reusable platform let the tenth agent ship in a fraction of the time the first one took. Without them, every new agent repeats the same access, review and validation work, and the cost curve never bends.

The critical path between stages

Each step up needs a different piece of work done first.

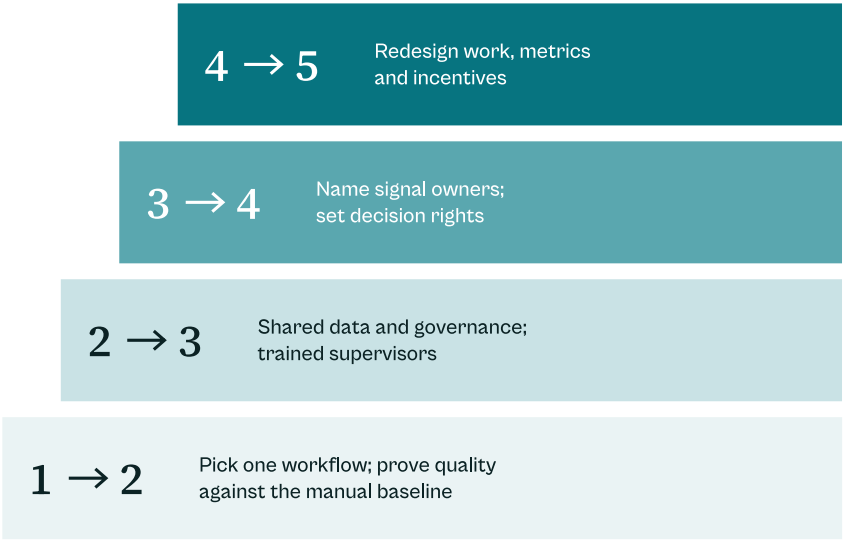

Moving up is an operating-model change, not a technology purchase.

THE COMMON TRAP

Companies often fund the visible part (the agent) and underfund the invisible part (data access, evaluation and governance). The agent demos well and then cannot be validated, trusted or reused. **Budget for foundations as part of the first workflow, not after it.**

The software-lite enterprise

Much business software exists to give a person a screen for one step in a workflow, such as an approval form, a tracking tool or a dashboard that shows whether something happened. When agents run the workflow, they work directly with the data and with permission to act.

Accellor expects these categories to shrink: single-workflow SaaS tools, departmental tools that move data between systems, dashboards built to prompt human action, custom internal apps that encode a process, and much of the integration software that connects them.

Systems of record remain. In life sciences, these are the CRM, the ERP, quality and regulatory systems, data infrastructure, and finance and human resources (HR) platforms. They hold the governed, auditable state of the business, and agents need a trusted place to read from and write to. As agents take over work that people did on screens, these platforms will be judged more on their APIs, data quality and permission models.

The target architecture has four layers, with security, identity, monitoring, evaluation and cost control running across all of them:

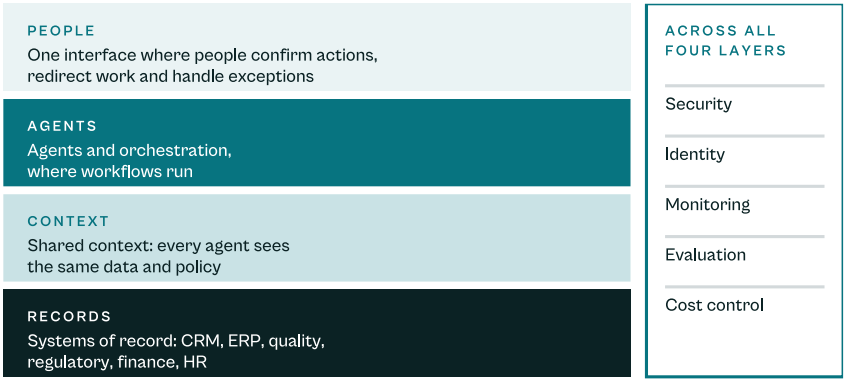

PART TWO

2

Life sciences, end to end

Four functions, two in pharma and two in medtech, tracked step by step through all five stages.

Stage and ceiling are different things. Moving up the stages changes who does the work. It does not lift the cap on a decision that affects patient safety, product quality or a regulatory submission.

Four examples, end to end

The four examples that follow track two pharma and two medtech functions through all five stages. Each example is one table, spread across two facing pages. The rows are the six main steps of the work, in order, with each step's consequence tier (page 15) and who owns the decision. The columns show what AI does at that step at each stage.

How to read the tables. Each cell names the most the AI does at that step. The shading deepens as the AI does more:

Person	Suggests	Drafts	Approval	Limits
No AI in this step. People do the work.	Answers questions and finds information when asked.	Produces work for a person to review.	Acts after a named person approves.	Acts inside set limits while a person monitors.

Each cell shows what is new at that stage. Work from earlier stages continues unless the cell says it is replaced. Read across a row to see one step mature, and down a column to see the whole function at one stage. A level can stay the same while the work grows: at Stage 4 the step draws on other functions' signals, and at Stage 5 the step itself is redesigned from outcomes on a governed cycle.

Tier 3 decisions never go past Drafts unless the workflow is validated, and the decision always stays with a named person. Where a Tier 3 step shows Limits, the agent only gathers, tracks or delivers information. It makes no decision.

Sampled review at Stage 5 means a person checks a defined sample of high-confidence drafts instead of every one, and still reviews every low-confidence or high-risk draft. It applies only to drafting, never to the decision itself, and only once the workflow is validated and the evidence supports it. Review returns to 100% after any model, data or process change, or a rise in errors.

The companies are illustrative composites, not client results.

Pharma example 1: Pharmacovigilance (PV)

A mid-size pharma company receives individual case reports from spontaneous reporting, literature, partners and patient support programs, and must meet strict reporting timelines. Case processors code each reported event using MedDRA, the standard medical dictionary for regulatory activities.

THE WORKFLOW, STAGES 1 TO 3

Step	1 Experiment	2 Embed	3 Delegate
1. Intake and triage Tier 1–3 · Case processor	Person	Person	**Drafts**: validity and duplicate checks; logging and routing run within limits
2. Data entry and coding Tier 3 · Case processor	**Suggests**: looks up terms	**Drafts**: extracts data, proposes MedDRA codes	**Drafts**: within the assembled case
3. Follow-up for missing information Tier 2 · Case processor	**Suggests**: wording	**Drafts**: request for missing fields (replaces wording help)	**Approval**: sends once approved
4. Narrative and medical assessment Tier 3 · Safety scientist	**Suggests**: summarizes articles	**Drafts**: full narrative	**Drafts**: proposed seriousness and expectedness
5. Expedited reporting Tier 3 · PV compliance lead	Person	Person	**Limits**: tracks reporting clocks; a person submits
6. Signal detection and communication Tier 3 · Safety physician	Person	Person	**Suggests**: surfaces candidate signals

Pharmacovigilance, continued

THE WORKFLOW, STAGES 4 AND 5

Step	4 Connect	5 Evolve
1. Intake and triage Tier 1–3 · Case processor	**Drafts**: matches duplicates across field, partner and patient-program sources	**Drafts**: rules retuned from reviewer corrections
2. Data entry and coding Tier 3 · Case processor	**Drafts**: codes match clinical and medical data	**Drafts**, sampled review: non-serious cases checked by sample
3. Follow-up for missing information Tier 2 · Case processor	**Approval**: adds questions prompted by open signals	**Limits**: pre-approved question sets only
4. Narrative and medical assessment Tier 3 · Safety scientist	**Drafts**: adds signal context	**Drafts**, sampled review: non-serious narratives only; every assessment decided by the scientist
5. Expedited reporting Tier 3 · PV compliance lead	**Limits**: adds partner and affiliate deadlines	**Limits**: forecasts workload from intake
6. Signal detection and communication Tier 3 · Safety physician	**Drafts**: signal assessment; validated signals delivered to named owners	**Drafts**: proposed updates to risk plans and protocols

What to notice

The beachhead is steps 2 and 4 (coding and narrative). They carry real judgment and teach the team to review AI output.

AI reaches every step by Stage 5, but the decisions never move. Steps 1, 2, 4 and 6 never go above Drafts. The Stage 5 gain is lighter review on non-serious cases, not more authority.

Only the Tier 2 step earns real authority: by Stage 5, follow-up requests go out within limits, using pre-approved question sets. Step 5 shows Limits only because the agent tracks clocks and deadlines; a person still submits every report.

Pharma example 2: Commercial field engagement

A brand team supports field representatives who call on physicians and other health care professionals (HCPs). Most of the work is Tier 2, but step 4 has Tier 3 moments: adverse events heard in the field, and medical questions reps must not answer themselves.

THE WORKFLOW, STAGES 1 TO 3

Step	1 Experiment	2 Embed	3 Delegate
1. Territory planning Tier 1 · Rep and district manager	**Suggests**: on request	**Suggests**: in approved tools, using CRM data	**Limits**: weekly priority list; rep can change it
2. Call preparation Tier 2 · Rep	**Suggests**: summarizes HCP publications	**Drafts**: pre-call plan with approved content	**Drafts**: as part of the weekly plan
3. The conversation Tier 2 · Rep	Person	Person	Person
4. Notes, adverse events and medical questions Tier 1–3 · Rep; safety and medical decide	Person	**Drafts**: call summary; flags possible adverse events	**Drafts**: adverse-event report to safety; CRM updated within limits
5. Follow-up with the HCP Tier 2 · Rep	**Suggests**: email wording	**Drafts**: follow-up from approved content (replaces wording help)	**Approval**: sends and books once the rep approves
6. Field insight Tier 1 · Brand lead	Person	Person	**Drafts**: weekly insight summary

Commercial field engagement, continued

THE WORKFLOW, STAGES 4 AND 5

Step	4 Connect	5 Evolve
1. Territory planning Tier 1 · Rep and district manager	**Limits**: adds medical and supply signals	**Limits**: planning logic redesigned from outcomes
2. Call preparation Tier 2 · Rep	**Drafts**: adds medical and safety updates	**Drafts**: plan redesigned from what reps used
3. The conversation Tier 2 · Rep	Person	Person
4. Notes, adverse events and medical questions Tier 1–3 · Rep; safety and medical decide	**Drafts**: direct links to safety and medical information	**Drafts**: detection tuned from safety's feedback
5. Follow-up with the HCP Tier 2 · Rep	**Approval**: timing informed by medical and marketing	**Limits**: sends approved resources an HCP requested
6. Field insight Tier 1 · Brand lead	**Limits**: tagged insight shared with medical and marketing	**Drafts**: proposed changes to content and call plans

What to notice

The beachhead is step 2 (call preparation). It is high-volume and needs judgment about the HCP.

The conversation stays with the rep at every stage. The agent's job is better preparation and more time for it.

AI never creates claims. It works only from reviewed content, so the review process stays the control point.

Step 4 is why commercial is not "all Tier 2." The agent helps catch adverse events and medical questions; safety and medical decide.

Medtech example 1: Manufacturing nonconformance and supply

A medtech company makes single-use surgical devices across three plants. When product does not meet specification, quality engineers investigate, a review board decides what happens to it and what corrective and preventive action (CAPA) is needed, and planners manage the effect on supply.

THE WORKFLOW, STAGES 1 TO 3

Step	1 Experiment	2 Embed	3 Delegate
1. Detect, log and contain Tier 3 · Quality engineer	Person	**Drafts**: classifies the defect, drafts the report	**Drafts**: containment plan; affected lots from genealogy
2. Gather records Tier 3 · No decision; engineer reviews	**Suggests**: searches operating procedures	**Suggests**: retrieves records on request	**Limits** (read-only): assembles the investigation pack
3. Root cause Tier 3 · Quality engineer	Person	Person	**Drafts**: likely causes with evidence
4. Disposition and CAPA Tier 3 · Review board; quality manager	Person	Person	**Drafts**: disposition and CAPA recommendations
5. Suppliers and production plan Tier 2 · Supplier quality; planner	**Suggests**: email wording	**Drafts**: supplier notice (replaces wording help)	**Approval**: expedites, revised schedule, alternate sources
6. Report and trend Tier 1 · Plant quality lead	**Suggests**: spreadsheet help	**Limits**: shift and nonconformance reports	**Limits**: tracks open investigations

Manufacturing nonconformance and supply, continued

THE WORKFLOW, STAGES 4 AND 5

Step	4 Connect	5 Evolve
1. Detect, log and contain Tier 3 · Quality engineer	**Drafts**: adds field failures and shipped product	**Drafts**, sampled review: known low-severity defects only; containment always approved
2. Gather records Tier 3 · No decision; engineer reviews	**Limits**: adds supplier and field data	**Limits**: pack redesigned around what investigators use
3. Root cause Tier 3 · Quality engineer	**Drafts**: uses cross-plant and field data	**Drafts**: draws on a growing library of confirmed causes
4. Disposition and CAPA Tier 3 · Review board; quality manager	**Drafts**: shows supply and customer impact	**Drafts**: logic updated from board decisions; the board decides
5. Suppliers and production plan Tier 2 · Supplier quality; planner	**Approval**: updates customer order promises	**Limits**: re-sequencing inside the approved plan
6. Report and trend Tier 1 · Plant quality lead	**Limits**: cross-plant trends shared	**Limits**: drafts redesign proposals for change control

What to notice

The beachhead is step 1 (classifying and drafting the report). It is high-volume and builds the labeled history later stages depend on.

Step 2 is the biggest time saving and the safest delegation: read-only assembly with no decision.

Disposition and CAPA stay with the board and the quality manager at every stage. Sampled review applies only to classifying known defects.

Medtech example 2: Field service

A medtech company services imaging systems installed in hospitals. Some service events are also complaints, and some of those must be reported to regulators, so field service has real Tier 3 exposure.

THE WORKFLOW, STAGES 1 TO 3

Step	1 Experiment	2 Embed	3 Delegate
1. Detect and diagnose Tier 1 · Remote service engineer	**Suggests**: searches manuals	**Suggests**: approved knowledge base	**Drafts**: diagnosis from history and logs; cases opened from alerts within limits
2. Plan the visit Tier 2 · Dispatcher	Person	Person	**Approval**: books the technician, orders parts
3. Repair on site Tier 2 · Technician	Person	Person, with **Suggests**: repair steps on request	Person, with **Suggests**: steps for the diagnosed fault
4. Service report and complaint screening Tier 2–3 · Technician; complaint handling	Person	**Drafts**: service report; flags possible complaints	**Drafts**: complaint file
5. Reportability Tier 3 · Quality	Person	Person	**Drafts**: assessment for quality
6. Close and learn Tier 1–2 · Service sales; service operations	**Suggests**: email wording	**Drafts**: customer summary (replaces wording help)	**Approval**: quotes and renewals; dashboards within limits

Field service, continued

THE WORKFLOW, STAGES 4 AND 5

Step	4 Connect	5 Evolve
1. Detect and diagnose Tier 1 · Remote service engineer	**Drafts**: adds failure data from all sites	**Drafts**: cases opened before failure; logic tuned from confirmed repairs
2. Plan the visit Tier 2 · Dispatcher	**Approval**: linked to parts planning	**Limits**: rebooks within the customer's agreed window
3. Repair on site Tier 2 · Technician	Person, with **Suggests**: fixes that worked elsewhere	Person, with **Suggests**: procedures updated from confirmed repairs
4. Service report and complaint screening Tier 2–3 · Technician; complaint handling	**Drafts**: checks against installed-base trends	**Drafts**: criteria retuned; every flag still reviewed
5. Reportability Tier 3 · Quality	**Drafts**: adds trend data	**Drafts**: guidance updated from past decisions
6. Close and learn Tier 1–2 · Service sales; service operations	**Approval**: renewals use installed-base data; failure trends shared with quality and design	**Approval**: contracts built around predicted uptime; drafts bulletin and training changes

What to notice

The beachhead is step 4 (report drafting and complaint screening). The key measure is complaints missed compared with manual review.

Complaints often first surface in the field, in steps 1 and 4. Regulatory reporting timelines start when the company first learns of a possible death or serious injury, wherever in the company that happens.

The repair stays with the technician at every stage. Complaint screening stays at full review, because a missed complaint is the costliest error.

PART THREE

3

Next Steps

Take our self-check quiz to see where your team is today and then create a plan to get more out of AI.

Using AI is one thing. Getting real value out of it is hard. Use our framework to take an honest look at where your team is today and then use our guide to build AI that *earns* its place in the organization.

Self-check

Think of one function and pick the answer closest to how its most important workflow runs today. Score A as 1, B as 2, C as 3, D as 4 and E as 5, then add up your score. A company is rarely at one stage everywhere, so repeat the check for each function.

Function: _____

1. How is AI used in this function today?

- ☐ **A** People try tools on their own
- ☐ **B** AI runs inside one workflow, and people review it
- ☐ **C** Agents carry multi-step work, and experts supervise
- ☐ **D** Agents share signal with other functions
- ☐ **E** Work is redesigned around what agents learn

2. Is there a list of every agent in use?

- ☐ **A** No
- ☐ **B** Each team keeps one
- ☐ **C** One list, with owners and permissions
- ☐ **D** One list, including the signals each agent shares
- ☐ **E** One list, with performance reviewed on a set cycle

3. How do agents reach CRM and ERP data?

- ☐ **A** They don't
- ☐ **B** Through exports and copies
- ☐ **C** Governed APIs, for this function's data
- ☐ **D** Governed APIs, shared across functions
- ☐ **E** Governed APIs, with data owners fixing what agents flag

4. When AI finishes a task, what happens next?

- ☐ **A** A person copies the output
- ☐ **B** A person reviews it and passes it on
- ☐ **C** An agent takes the next step, and an expert supervises
- ☐ **D** The result reaches other functions as a governed signal
- ☐ **E** The outcome feeds back into how the work is designed

5. How does a finding in this function reach the functions it affects?

- ☐ **A** By email and meetings, if at all
- ☐ **B** In periodic reports
- ☐ **C** Agents pass it on within this function only
- ☐ **D** As a governed signal with a named owner
- ☐ **E** As a governed signal that also updates processes and plans

6. Are there written rules for what agents may do without approval?

- ☐ **A** No
- ☐ **B** For one workflow
- ☐ **C** For every agent, by consequence tier
- ☐ **D** Yes, including what agents may share across functions
- ☐ **E** Yes, revised on a set cycle as evidence comes in

7. Is there one record for each customer, product lot and location?

- ☐ **A** No
- ☐ **B** In some systems
- ☐ **C** Yes, in the core systems
- ☐ **D** Yes, shared across functions
- ☐ **E** Yes, updated as events happen

8. How is the value of AI measured?

- ☐ **A** It isn't
- ☐ **B** By task, against the manual baseline
- ☐ **C** By workflow, including exceptions and time to decision
- ☐ **D** By the time from a signal to a decision across functions
- ☐ **E** By business outcome, with work redesigned from the results

9. Has any software been retired because agents replaced it?

- ☐ **A** No
- ☐ **B** Under discussion
- ☐ **C** One or two tools
- ☐ **D** Several tools, across functions
- ☐ **E** Retirement is planned every year

10. Is every agent action logged, attributable and reversible where possible?

- ☐ **A** No
- ☐ **B** For one workflow
- ☐ **C** For every agent
- ☐ **D** For every agent and every signal shared
- ☐ **E** For all of them, with regular reviews that update the rules

YOUR SCORE

10–17	18–25	26–33	34–41	42–50
Stage 1 **Experiment**	Stage 2 **Embed**	Stage 3 **Delegate**	Stage 4 **Connect**	Stage 5 **Evolve**

Your score: _____ Circle where you landed.

Use your score with the 90-day plan on the facing page.

Your first 90 days

Moving up one stage starts with one workflow. Your 90-day plan can be run in six steps:

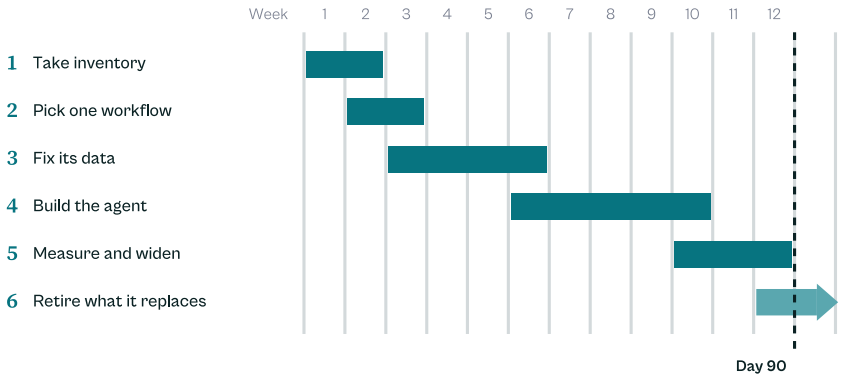

1 **Take inventory.** Weeks 1–2. List the agents, AI tools, data sources and workflows in use, and where they overlap.

2 **Pick one workflow.** Weeks 2–3. Choose one with clear value and data that is close to ready, such as case coding, call preparation, defect reports or complaint screening.

3 **Fix its data.** Weeks 3–6. Give the workflow's key records an owner, one authoritative source and API access.

4 **Build a propose-and-confirm agent.** Weeks 6–10. The agent gathers the inputs, prepares the action and asks a named person to approve it.

5 **Measure and widen.** Weeks 10–12. Track accuracy and approval rates. Where the tier allows, move proven actions from Approval to Limits: the agent acts inside set limits and reports what it did.

6 **Retire what it replaces.** Week 12 onward. Decide which tool, report or handoff the workflow no longer needs.

TALK TO US

Bring your self-check score to a 30-minute conversation with Accellor's Healthcare and Life Sciences team. We will place each of your functions on the five stages, show where autonomy is capped, and help you pick the one workflow to start with.

accellor.com/what-we-do/industries/life-sciences

Accellor has worked with healthcare and life sciences organizations for more than eight years and on more than 70 projects. It is a Salesforce Summit Partner and a Microsoft Solutions Partner for Data & AI, with teams in the United States, the United Kingdom and India.

Sources and notes

1. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," press release, June 25, 2025.

2. McKinsey & Company, "The State of AI," https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai.

3. Boston Consulting Group, "The Leader's Guide to Transforming with AI," https://www.bcg.com/featured-insights/the-leaders-guide-to-transforming-with-ai.

4. Deloitte, "State of AI in the Life Sciences & Health Care Industry: The Untapped Edge," March 2026, https://www.deloitte.com/content/dam/assets-shared/docs/about/2026/state-of-ai-lshc.pdf.

5. Salesforce, "Multi-Agent Adoption to Surge 67% by 2027 as Enterprises Race Toward Agentic Transformation – Unified Architecture Key to Success," announcement of the 2026 Connectivity Benchmark Report, February 5, 2026, https://www.salesforce.com/news/stories/connectivity-report-announcement-2026/.

6. Zylo, "Zylo's 2026 SaaS Management Index Finds AI-Native App Adoption Is Surging, with ChatGPT Now the Most Expensed App," January 29, 2026, https://zylo.com/news/2026-saas-management-index.

7. U.S. Food and Drug Administration, "Computer Software Assurance for Production and Quality Management System Software," final guidance, February 2026, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/computer-software-assurance-production-and-quality-management-system-software.

The four examples in Part Two are illustrative composites, not client results. All sources were accessed on October 6, 2026.

© 2026 Accellor. All rights reserved.

AI that Earns its Place

A field guide to the five stages of AI maturity

Life Sciences Edition

accellor