CatalystDx Executive Education Series

Paper No. 1

The Agentic Enterprise

From Generative AI to Autonomous Organizational Workflows

Peter M. Serpentino

CatalystDx

EXECUTIVE PREMISE
The competitive question is shifting from "How do we give employees AI?" to "Which workflows should AI agents operate, under what decision rights, and with what evidence of control?"

About this series

CatalystDx Executive Education translates fast-moving advances in artificial intelligence, R&D strategy and organizational design into practical operating frameworks for leaders in diagnostics, medical technology and the broader life-sciences enterprise. The series is designed for boards, CEOs, CTOs, R&D leaders, strategy executives and transformation teams that must convert AI capability into measurable organizational performance.

How to use this paper

Read this paper as an executive orientation, not a software implementation manual. Its purpose is to establish a common language for agentic AI, define a maturity path, show what an enterprise architecture can look like, and provide a disciplined way to identify where autonomy creates value—and where human authority must remain explicit.

CATALYSTDX POINT OF VIEW
Agentic AI should be treated as organizational redesign enabled by software. The durable value will come less from deploying isolated agents and more from redesigning workflows, decision rights, data access, evaluation and governance around a human–agent operating model.

Contents

1. Executive Summary
2. Why Agentic AI Matters Now
3. From Generative AI to Agentic Work
4. Anatomy of an Enterprise Agent
5. Architecture of the Agentic Enterprise
6. The CatalystDx Maturity Model
7. Diagnostics and Life-Sciences Applications
8. From Task Automation to Workflow Transformation
9. Executive Implications
10. Governance, Risk and Decision Rights
11. 90-Day Executive Action Agenda
12. Agentic Enterprise Assessment Checklist
13. Conclusion
14. Preview: Paper No. 2 — From Copilots to Agents

References

1. Executive Summary

Generative AI introduced an extraordinarily capable new interface to knowledge work. Agentic AI changes the unit of value. Instead of asking a model to produce an answer, a user can increasingly assign a goal to software that plans, uses tools, retrieves enterprise data, takes actions, checks results and continues until a defined outcome is reached.

That difference is strategic. A generative model can draft a regulatory summary. An agent can monitor new guidance, determine which programs may be affected, retrieve the relevant design history and regulatory plans, prepare a gap analysis, route it to the accountable owner and track closure. One produces content. The other begins to operate a workflow.

The enterprise opportunity therefore should not be framed as “deploy more AI assistants.” It should be framed as a redesign of work: which decisions remain human, which tasks can be delegated, which workflows can be orchestrated, what data and tools an agent may access, how performance is evaluated, and how errors are detected and recovered.

- Agentic AI is best understood as a goal-directed software operating model, not simply a more conversational interface.
- The highest-value use cases typically cross several steps or systems and require persistent context, tool use, exception handling and measurable completion criteria.
- Human–agent teams are likely to be the dominant near-term design. Current enterprise research finds leaders expect substantial process redesign around agents, while readiness for cross-functional orchestration remains low.[1]
- In diagnostics and other regulated life-sciences settings, autonomy must be matched to consequence. Agents can prepare evidence, detect gaps and coordinate work; accountable humans should retain high-consequence scientific, quality, regulatory, clinical and release decisions.
- The maturity challenge is organizational. Data quality, identity, system access, workflow ownership, evaluation, cybersecurity and governance are prerequisites to scale.
- The executive starting point is not “Which agent platform should we buy?” It is “Which recurring workflows are valuable, measurable, sufficiently standardized and safe enough to redesign?”

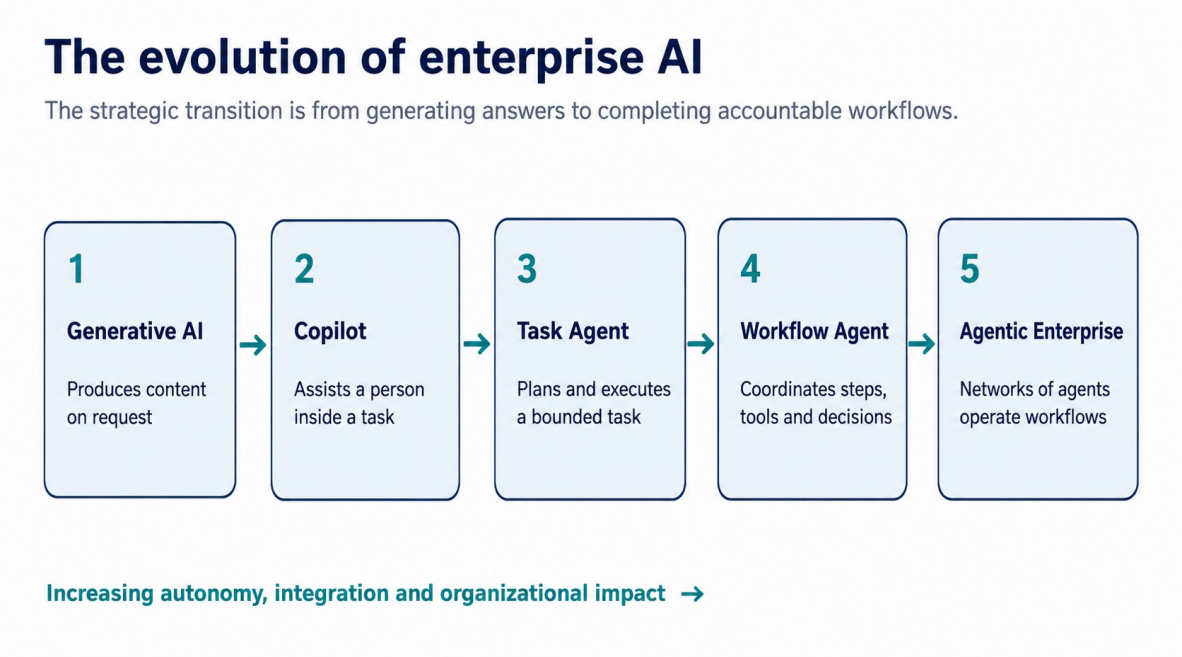

Exhibit 1. The evolution from generative assistance to an agentic operating model. Source: CatalystDx framework.

2. Why Agentic AI Matters Now

The agentic shift is arriving at the same time that many enterprises are confronting a structural limit in knowledge work: the volume of information, coordination and digital interaction is increasing faster than teams can absorb it. Microsoft's 2026 Work Trend Index frames the leadership task as rearchitecting work as agents take on more execution and humans retain greater agency over direction and outcomes.[2] This is a useful distinction. The value proposition is not "AI replaces people." It is that scarce human attention moves away from routine execution and toward judgment, prioritization, relationship-building and accountability.

The organizational implications are already visible. Deloitte reported in August 2026 that 74% of surveyed leaders expect nearly half of business processes to be redesigned or rebuilt around AI agents within four years; at the same time, only 5% said their processes were highly prepared for agents and only 15% reported scaled, orchestrated cross-functional multi-agent adoption.[1] Expectations are running ahead of operating readiness.

That gap creates both opportunity and risk. The winners will not necessarily be the organizations that experiment first. They will be the organizations that learn fastest how to redesign repeatable work around reliable human–agent systems—and that build governance quickly enough to scale what works.

Executive shift	Generative-AI era	Agentic-enterprise era
Primary unit of value	Prompt / response	Workflow / outcome
Human role	Ask, edit, verify	Set goals, approve exceptions, own outcomes

Executive shift	Generative-AI era	Agentic-enterprise era
AI role	Generate or analyze	Plan, use tools, act, observe, evaluate
Integration	Often standalone	Embedded in enterprise systems and processes
Success metric	Output quality and productivity	Cycle time, completion, economics, control and outcome quality
Governance focus	Content and data risks	Authority, identity, action, monitoring, recoverability and auditability

What Is Genuinely New?

Autonomous software is not new. Rule engines, expert systems, robotic process automation, workflow management systems, and classical software agents have existed for decades. What is genuinely new is the emergence of **large language models (LLMs)** as a general-purpose interface between human intent, enterprise knowledge, software tools, and increasingly, organizational workflows.

Large Language Models: A New Interface to Enterprise Knowledge and Work

A **large language model**, or LLM, is a form of artificial intelligence trained on extremely large collections of text, code, documents, and other information. At its foundation, an LLM learns statistical relationships among words, concepts, patterns, and representations. Given a question, instruction, document, dataset, or other context, it can generate a response by predicting and constructing an appropriate sequence of information.

That technical description, however, understates its importance to the enterprise.

Traditional software generally requires people to translate a business problem into a predefined set of fields, rules, menus, database queries, or programmed instructions. An LLM can instead accept much of the information in the form in which organizations naturally create it: reports, emails, customer comments, research papers, meeting notes, requirements documents, experimental results, regulatory guidance, competitive intelligence, technical specifications, and management instructions.

The LLM therefore becomes a **reasoning and language interface across otherwise fragmented organizational information**.

A properly configured enterprise system built around an LLM can, for example:

- read thousands of customer comments and identify recurring unmet needs;
- compare customer requirements with competitive products and market trends;
- summarize scientific and technical literature relevant to a development program;

- examine experimental results against predefined technical milestones;
- identify inconsistencies among product requirements, verification plans, and risk files;
- prepare a Stage-Gate evidence package from information residing across multiple functions;
- monitor launch-readiness metrics and identify unresolved risks; and
- communicate different versions of the same underlying analysis to scientists, executives, regulators, commercial teams, or customers.

The significance is not simply that the computer can generate better text. The significance is that **a common AI layer can interpret, connect, synthesize, and act upon information that previously required extensive human translation between organizational silos.**

The LLM Is Not the Agent

For executives, an important distinction is necessary.

An LLM by itself is not an autonomous enterprise agent. It is better understood as a powerful cognitive engine that can interpret language, synthesize evidence, generate alternatives, and determine possible next steps.

It becomes part of an **agentic system** when additional capabilities are placed around it:

Objective → LLM interpretation and planning → Enterprise data → Approved tools and APIs → Action → Observation of results → Evaluation → Next action or escalation

The surrounding architecture can provide the agent with access to controlled data, search systems, databases, laboratory information, project-management systems, CRM platforms, regulatory repositories, financial models, or other enterprise applications.

Guardrails can then define what the agent may read, what it may recommend, what it may change, what requires verification, and what must always remain a human decision.

This distinction is particularly important in regulated industries. The objective should not be unconstrained autonomy. The objective should be **bounded, accountable autonomy within a defined operating process**.

Why Prescriptive Innovation Processes Are Particularly Well Suited to Agentic AI

Paradoxically, organizations with disciplined innovation processes may have an advantage in deploying agentic AI.

Agents perform best when the organization can clearly specify:

What is the objective?
What information is required?
What constitutes acceptable evidence?
What milestone must be achieved?

Who owns the decision?
What actions are permitted?
What conditions require escalation?
What defines successful completion?

A rigorous innovation-management system already answers many of these questions.

This makes established processes such as **market surveillance, Voice of the Customer, technology milestone management, product-development Stage-Gate, design controls, and product-launch readiness** natural environments for carefully governed agentic AI.

The opportunity is not to discard those processes. It is to make them **continuous, connected, evidence-aware, and increasingly intelligent.**

From Periodic Reviews to Continuous Innovation Intelligence

Consider the traditional innovation lifecycle:

Market Surveillance & Voice of the Customer
→ Opportunity Definition
→ Technology Feasibility
→ Product Development
→ Verification & Validation
→ Regulatory / Manufacturing Readiness
→ Product Launch
→ Post-Launch Surveillance

These processes are typically highly structured, but the information supporting them is often fragmented across people, documents, databases, functions, and meetings.

LLMs and agents can create a continuous intelligence layer across the entire process.

Market Surveillance and Voice of the Customer

Instead of relying primarily on periodic market studies, an AI-enabled process can continuously examine customer feedback, support records, competitive activity, publications, conference proceedings, patents, regulatory clearances, clinical trends, and commercial information.

A **Market and VOC Agent** could identify emerging needs, cluster customer problems, distinguish isolated requests from recurring patterns, and maintain an evolving evidence base supporting the product opportunity.

Human leaders still determine strategy. The agent dramatically increases the amount and frequency of evidence available to them.

Technology Milestone Process

During technology feasibility, an agent can monitor experimental results against predefined milestones and acceptance criteria.

Rather than asking only whether an experiment "worked," the system can continuously evaluate questions such as:

- Has analytical sensitivity reached the target?
- Is reproducibility improving?
- Are unresolved technical risks increasing or decreasing?
- Which failure modes recur across experiments?
- Is the technology trajectory sufficient to support progression?
- Which additional experiments would provide the greatest reduction in uncertainty?

The technology team retains scientific accountability, but management gains a more continuous and objective view of **technology readiness**.

Product Development Stage-Gate

Stage-Gate processes are especially amenable to agentic augmentation because each gate has an explicit purpose, required evidence, deliverables, risks, and decision authority.

A **Stage-Gate Agent** could continuously assemble evidence from R&D, systems engineering, clinical, regulatory, quality, manufacturing, marketing, and finance; compare the evidence against gate criteria; identify missing deliverables or contradictory assumptions; and prepare the decision package before the formal review.

The gate committee would not be replaced.

Its decision could become better informed.

Instead of discovering deficiencies during the meeting, management could know weeks earlier that manufacturing readiness is lagging technology readiness, that the clinical plan no longer supports the proposed claim, or that the commercial assumptions have changed since project authorization.

Agentic AI therefore creates the possibility of moving from **periodic Stage-Gate inspection to continuous Stage-Gate readiness management**.

Product Launch

The same approach extends through product launch.

A Launch Readiness Agent could monitor manufacturing capacity, supply-chain risks, regulatory status, training, service readiness, inventory, commercial materials, pricing, customer-site preparation, and early market signals.

It could continuously determine whether the organization is converging toward launch readiness and highlight dependencies that threaten the launch date or commercial outcome.

Again, the purpose is not to allow an AI system to authorize product release. The purpose is to ensure that executives making that decision have a more complete, current, and internally consistent body of evidence.

The Strategic Shift: From AI as a Tool to AI as an Operating Layer

This is the more consequential change.

Generative AI is commonly introduced as an individual productivity tool: write this report, summarize this document, analyze this spreadsheet, or answer this question.

Agentic AI creates a different possibility.

It allows intelligence to be embedded **inside the enterprise process itself**.

The progression is therefore:

AI assists an individual
→ AI assists a task
→ AI coordinates a workflow
→ AI monitors an end-to-end process
→ Networks of agents support the operating system of the enterprise

For innovation organizations, that could eventually mean that market evidence continuously informs product requirements; technical results continuously update technology-readiness assessments; changes in design automatically trigger downstream risk reviews; commercial assumptions remain connected to product decisions; and executives receive an ongoing assessment of whether a program is truly ready for the next commitment of resources.

That does not make disciplined innovation management obsolete.

It makes discipline more important.

The more precisely an organization defines its milestones, acceptance criteria, evidence requirements, decision rights, and governance, the more effectively it can deploy AI agents within those boundaries.

The central opportunity is therefore not to replace the innovation process with AI. It is to transform the innovation process into an intelligent, continuously learning, evidence-driven operating system.

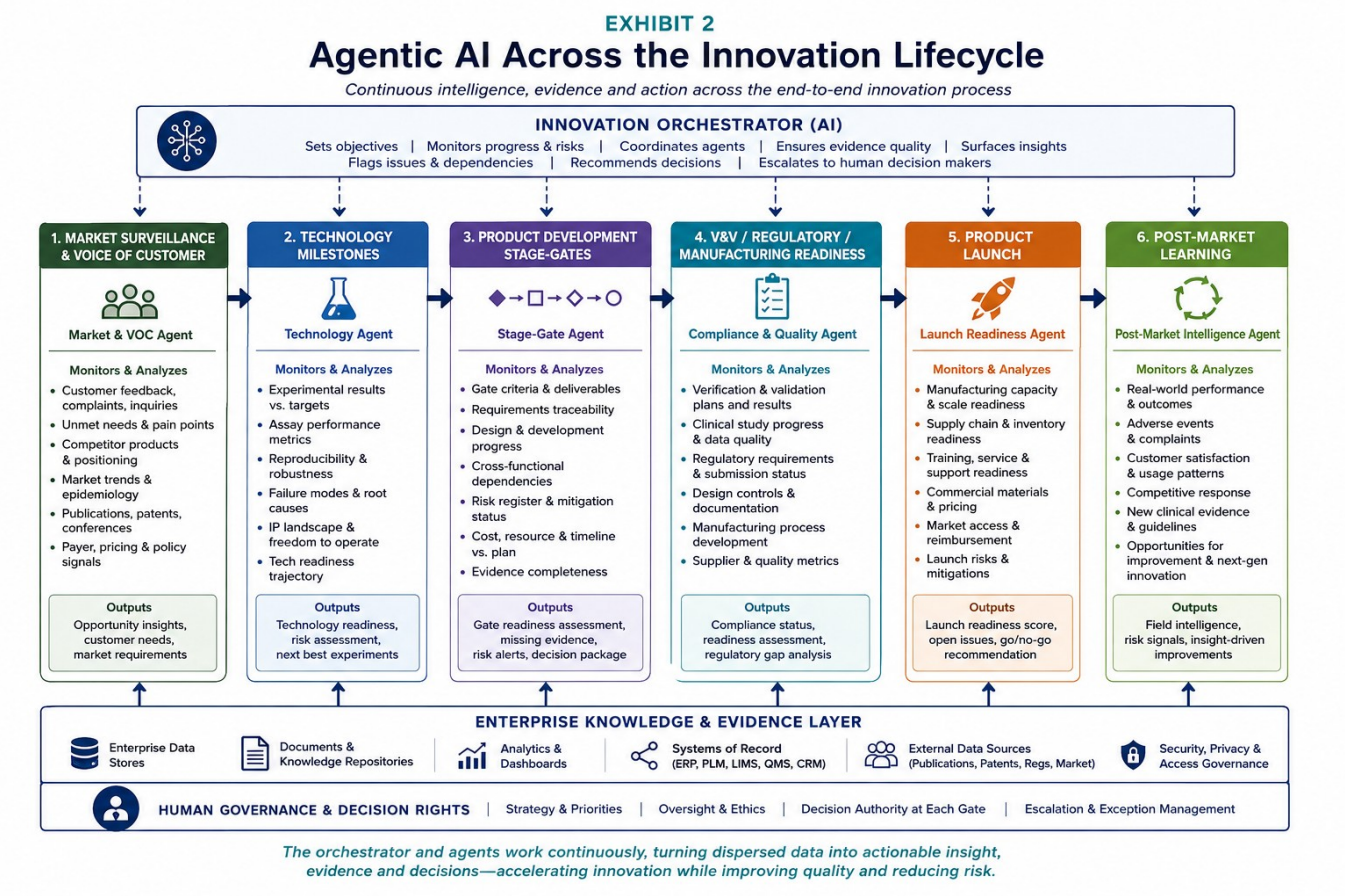

3. From Generative AI to Agentic Work

Executives need a vocabulary that distinguishes genuinely agentic behavior from marketing labels. CatalystDx recommends five levels of capability. The boundaries are not perfect, but they force a useful discussion about autonomy, integration and accountability.

Capability	What it does	Typical example	Primary limitation
Generative AI	Produces content or analysis in response to a prompt.	Summarize a clinical paper.	No persistent responsibility for completing work.
Copilot	Assists a person within an application or role.	Draft a design-review summary from selected documents.	Human initiates and sequences the work.
Task agent	Pursues a bounded objective using tools and rules.	Reconcile open CAPA actions and prepare an exception list.	Narrow scope; limited cross-functional coordination.
Workflow agent	Coordinates multiple steps, systems and checkpoints.	Prepare a monthly R&D gate-readiness packet across functions.	Requires process integration, monitoring and governance.

Capability	What it does	Typical example	Primary limitation
Agentic operating model	Networks of governed agents execute recurring workflows under human decision rights.	Continuous program monitoring, evidence collection and escalation across an R&D portfolio.	Organizational complexity, trust, security and control become central.

Three tests for “agentic”

1. Goal persistence: Can the system remain responsible for an objective across multiple steps rather than respond once?
2. Environmental action: Can it use approved tools, data or systems to change the state of the work—not merely describe what someone else should do?
3. Closed-loop adaptation: Can it observe the result, evaluate progress and choose the next step or escalate an exception?

A system does not need unlimited autonomy to be agentic. In most enterprises, bounded autonomy is preferable. A highly useful agent may be allowed to search, analyze, prepare, reconcile and route work while being prohibited from making financial commitments, changing controlled records, releasing product or making regulated decisions without explicit approval.

EXECUTIVE QUESTION
If the AI stopped after producing an answer, would the workflow still need a person to perform most of the remaining coordination? If yes, the organization is probably still in the copilot stage—not yet redesigning the workflow.

4. Anatomy of an Enterprise Agent

An enterprise agent can be modeled as a closed loop. The model interprets the objective and plans; tools give it the ability to retrieve data or perform actions; memory retains relevant context; evaluation determines whether the work meets quality or policy thresholds; guardrails constrain what it may do; and human escalation handles ambiguity or consequences beyond its authority.

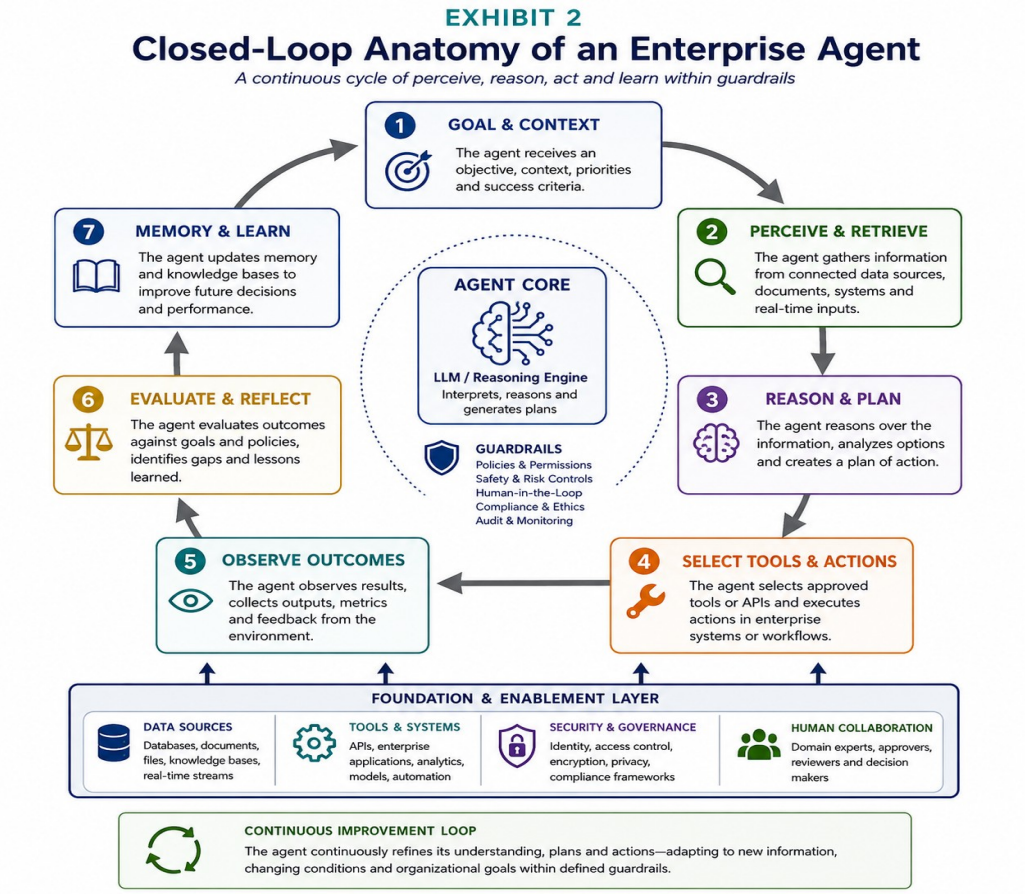

Exhibit 2. Closed-loop anatomy of an enterprise agent. Source: CatalystDx framework.

The eight design elements executives should require

Element	Executive meaning
Goal	A clearly bounded business outcome, not an open-ended instruction to “be helpful.”
Context	The records, policies, history and environment required to interpret the goal correctly.
Model / reasoning	The capability used to interpret evidence, plan and select actions.
Tools	Approved functions such as search, database queries, calculations, workflow updates or communications.
Memory	Durable context relevant to the task, with controls on what may be stored and reused.

Element	Executive meaning
Evaluation	Tests for factuality, completeness, policy compliance, business rules and outcome quality.
Guardrails	Permissions, prohibited actions, spending or data limits, segregation of duties and escalation thresholds.
Observability	Logs, traces, performance metrics and audit evidence that allow the organization to reconstruct what happened.

This architecture creates an important management implication: agent performance is a systems property. A capable model cannot compensate indefinitely for poor source data, ambiguous process ownership, excessive permissions or a workflow with no measurable definition of "done." Conversely, a well-designed workflow can often create value with moderate autonomy because the system knows what evidence to gather, when to stop and when to ask a person.

5. Architecture of the Agentic Enterprise

At scale, enterprises will not manage a collection of disconnected "bots." They will need an operating architecture that connects executive intent, workflow orchestration, specialized agents, enterprise systems, data, identity and governance. The architecture below is deliberately technology-agnostic. Its purpose is to show the management layers that must exist regardless of which software vendors are selected.

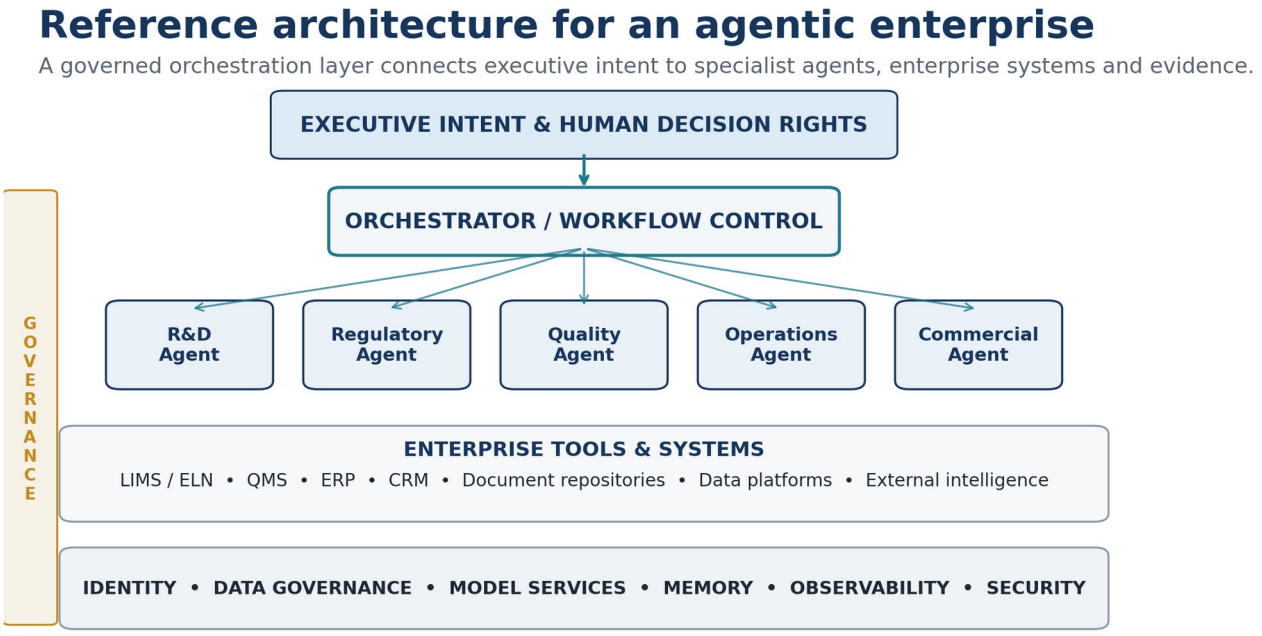

Exhibit 3. Reference architecture for an agentic enterprise. Source: CatalystDx framework.

Layer 1 — Executive intent and decision rights

Humans remain accountable for objectives, priorities, budgets, risk tolerance and high-consequence decisions. An agent should not inherit authority merely because a software connector makes an action technically possible.

Layer 2 — Orchestration and workflow control

An orchestrator converts an objective into coordinated work: assigning tasks, maintaining state, resolving dependencies, evaluating results and escalating exceptions. For complex processes, this layer is more strategically important than any single specialist agent.

Layer 3 — Specialized agents

Specialized agents can combine role-specific instructions, data access, tools and evaluation criteria. A Regulatory Agent should operate differently from an R&D Experiment Agent because the evidence, acceptable uncertainty, permissions and human owners are different.

Layer 4 — Enterprise systems and evidence

Agents create value when they can work with the systems where the organization's truth resides: LIMS, ELN, QMS, ERP, CRM, requirements systems, document repositories, data platforms and external intelligence. Read-only retrieval is generally the lowest-risk starting point. Write access, controlled-record changes and external actions require progressively stronger controls.

Layer 5 — Governance infrastructure

Identity, access, model services, data governance, security, memory, observability and evaluation form the foundation. As autonomy rises, these capabilities become operating infrastructure rather than optional AI features.

6. The CatalystDx Agentic Enterprise Maturity Model

Many organizations describe themselves as advanced in AI because a large percentage of employees have access to generative tools. That is not the same as agentic maturity. CatalystDx defines maturity by the degree to which repeatable work has been redesigned around governed autonomy and measurable outcomes.

EXHIBIT 4

CatalystDx Agentic Enterprise Maturity Model

A progression from isolated AI tools to an autonomous, continuously learning enterprise

DIMENSION	0 Ad Hoc / No AI	1 AI Assisted (Tools)	2 Task Automation (Agents)	3 Workflow Orchestration	4 Enterprise Integration	5 Autonomous Enterprise
Description	No formal AI use. Manual, siloed processes.	Employees use AI tools individually to improve productivity.	Agents automate specific tasks with defined inputs and outputs.	Agents coordinate multi-step workflows across functions and systems.	Agents integrated enterprise-wide with shared data, standards and governance.	Adaptive, goal-driven enterprise where agents operate autonomously within defined guardrails.
AI Role	None	Information retrieval, content generation, analysis assistance	Execute defined tasks, follow rules, trigger actions	Coordinate workflows, manage dependencies, route exceptions	Optimize end-to-end processes, allocate resources, predict and recommend	Continuously optimize, learn and adapt; initiate actions to achieve goals
Human Involvement	All manual	Human performs most work; AI used as a helper	Humans set direction and review outputs	Humans manage by exception; oversee workflows	Humans set strategy and policies; limited oversight	Humans define strategy and guardrails; monitor outcomes
Integration & Data	Siloed data, limited sharing	Limited integrations; ad hoc data connections	Connected to select systems and data sources	Integrated across multiple systems and data domains	Enterprise data fabric; real-time, trusted, governed	Unified semantic layer; real-time, self-improving data
Outcomes	Low efficiency, high variability, limited insight	Improved individual productivity and faster information access	Lower cycle time, fewer errors, consistent task execution	End-to-end visibility, better coordination, faster decisions	Optimized performance, predictive insights, higher quality and compliance	Autonomous operation, continuous learning, maximized value creation
Example State	Spreadsheets, email, manual reports	ChatGPT for drafting, copilots, standalone AI tools	RPA + AI for reporting, data extraction, alerts	Multi-agent workflows across R&D, quality, regulatory, supply chain	Connected enterprise agents across functions and value chain	Self-optimizing value networks with autonomous agents

INCREASING MATURITY

Key Takeaway: Maturity is not about technology alone—it is about data, governance, people, processes and culture. Progress in each dimension builds the foundation for enterprise-wide agentic AI success.

Exhibit 4. CatalystDx Agentic Enterprise Maturity Model. Source: CatalystDx framework.

A useful interpretation of the levels

- Levels 0–1 are primarily about individual productivity.
- Level 2 embeds assistance into roles and applications but leaves process choreography largely human.
- Level 3 introduces bounded delegation: the agent owns a task within explicit permissions and completion criteria.
- Level 4 is the major organizational transition because the unit of design becomes the end-to-end workflow rather than the individual task.
- Level 5 does not mean “fully autonomous company.” It means the enterprise deliberately assigns recurring execution to governed agent systems while humans retain strategic, fiduciary and high-consequence decision rights.

Most organizations should not attempt to “skip” directly to Level 5. The capabilities required for safe scaling—identity, data access, process standardization, evaluation, security, change management and incident handling—are learned through bounded deployments. The objective is not maximal autonomy. It is the economically optimal level of autonomy for each workflow.

7. Diagnostics and Life-Sciences Applications

Diagnostics and life sciences are compelling proving grounds for agentic AI because the work is knowledge-intensive, cross-functional, document-heavy and evidence-driven. They are also

demanding environments: quality systems, design controls, clinical evidence, cybersecurity and regulatory obligations mean that speed without traceability can destroy value.

Recent scientific literature points toward increasingly agentic research environments. A 2025 Nature Machine Intelligence commentary described “agentic science” as the use of semi-autonomous agents capable of reasoning, planning and interacting with digital and physical environments.[5] In a 2025 Nature study, a multi-agent “Virtual Lab” was used to design SARS-CoV-2 nanobodies, illustrating how specialized AI roles can collaborate on open-ended scientific work.[6] These examples do not imply that autonomous agents are ready to own regulated product decisions. They demonstrate that the frontier is moving from question answering toward coordinated scientific problem solving.

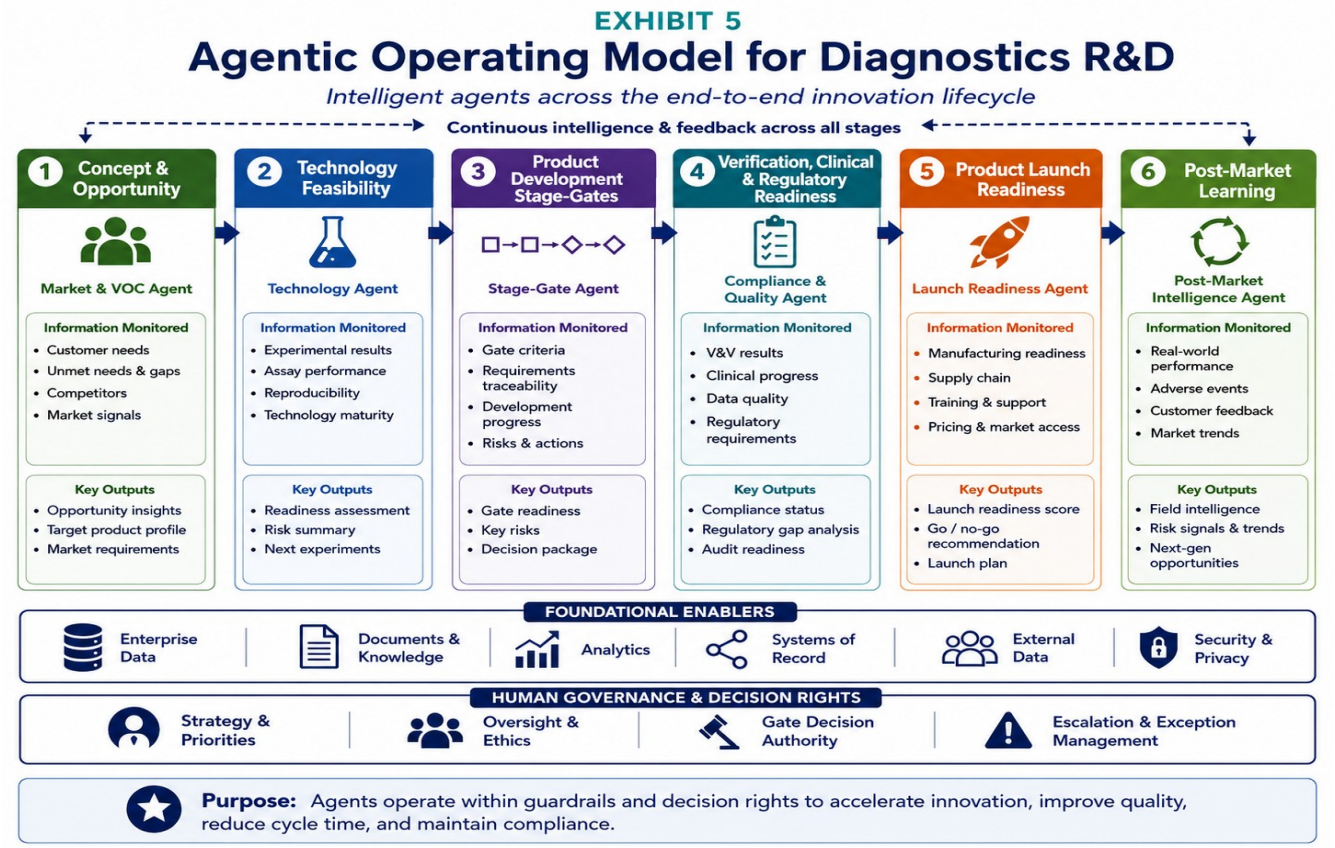

Exhibit 5. Illustrative agentic operating model for diagnostics R&D. Source: CatalystDx framework.

High-value diagnostics use cases

Workflow	Illustrative agent role	Human authority that remains explicit
Scientific intelligence	Continuously monitor literature, patents, competitors and emerging biomarkers; summarize implications for active programs.	Portfolio priorities; interpretation of strategic significance.
Assay development	Compare experimental results with historical data, identify anomalous trends and propose next-best experiments.	Scientific hypothesis selection; protocol approval; acceptance of evidence.
Requirements / systems engineering	Trace requirements to verification evidence, detect missing links and prepare review packets.	Requirements approval and design-review decisions.
Clinical development	Draft study documents, monitor operational exceptions and reconcile evidence packages.	Protocol approval; clinical interpretation; safety decisions.
Regulatory intelligence	Monitor guidance and policy changes; identify programs potentially affected; prepare gap analyses.	Regulatory strategy and commitments to authorities.
Quality systems	Triage recurring deviations, retrieve precedent, draft investigations and flag overdue actions.	Root-cause acceptance, CAPA approval, disposition and release decisions.

Workflow	Illustrative agent role	Human authority that remains explicit
Program management	Assemble cross-functional status, risk, readiness and decision evidence for stage gates.	Gate decisions, resource allocation and capital release.

Deloitte has similarly identified medtech opportunities in R&D documentation, regulatory monitoring, shared services, commercial content and supply-chain workflows.[7] The strategic opportunity for diagnostics companies is therefore broader than embedding AI into the product itself. Agentic AI can change how the company develops, validates, manufactures, supports and commercializes the product.

8. From Task Automation to Workflow Transformation

The most common failure mode in enterprise AI is to automate a task without redesigning the process around it. Drafting a report 80% faster produces limited economic value if people still spend days searching for source evidence, reconciling versions, emailing reviewers, checking missing actions and rebuilding status manually.

Agentic transformation starts with workflow decomposition. The executive team should map how work actually moves—not how the procedure says it moves—then decide where intelligence, coordination and authority belong.

Illustrative case: the R&D gate-readiness process

Traditional workflow	Agentic redesign
Program manager requests status updates from functions.	Program Orchestrator continuously reads approved milestone, risk and evidence sources.
Functions manually assemble slides and spreadsheets.	Specialist agents prepare standardized evidence summaries and exception lists.
Readiness gaps are discovered late in the review cycle.	Rules and evaluations flag missing verification, manufacturing, regulatory or market evidence as it emerges.
Meeting time is spent reconciling facts.	Meeting packet presents reconciled evidence, uncertainty and unresolved exceptions.
Committee decides whether to proceed.	Committee retains decision authority; the agent records rationale and routes actions.

CATALYSTDX DESIGN PRINCIPLE

Automate evidence preparation before automating judgment. In high-consequence workflows, the first major economic gain often comes from ensuring that decision-makers receive complete, current and traceable evidence—not from delegating the final decision to AI.

The workflow selection test

A workflow is a strong candidate for agentic redesign when most of the following are true:

- It recurs frequently enough that cycle-time or labor savings accumulate.
- The desired outcome can be defined and measured.
- The work requires coordination across multiple information sources or systems.
- Exceptions can be categorized and routed to accountable humans.
- The agent can operate initially with limited permissions or reversible actions.
- The required evidence is sufficiently accessible, structured or retrievable.
- There is an identifiable process owner who can define policy and accept the outcome.

Workflows with ambiguous ownership, poor data provenance, highly variable goals or irreversible high-consequence actions are usually poor candidates for early autonomy. They may still benefit from AI assistance, but they should not be forced prematurely into agentic operation.

9. Executive Implications: The Organization Changes Around the Agent

The largest implications of agentic AI are managerial, not technical. When software can pursue goals across systems, traditional assumptions about roles, spans of control, workflow ownership, metrics and decision rights need to be reconsidered.

1. Jobs become portfolios of human and agent work

A role should no longer be treated as an indivisible job description. Leaders can decompose it into activities: judgment, relationship, creation, analysis, coordination, transaction, monitoring and approval. Some activities remain human-first; others become agent-first; many become hybrid. This decomposition is more useful than debating whether “AI will replace the job.”

2. Management shifts from supervising activity to governing outcomes

When an agent can perform ten steps without constant human prompting, managers need visibility into goals, quality thresholds, exceptions and results rather than every keystroke. Observability and evaluation become management systems.

3. Process ownership becomes more important

An agent cannot compensate for a workflow that has no accountable owner. Someone must define what good looks like, what data are authoritative, which exceptions matter, what actions are permitted and who accepts the output.

4. Functional silos become a design constraint

The highest-value workflows often cross R&D, quality, regulatory, manufacturing and commercial functions. If each function builds isolated agents, the enterprise may simply automate silos. Orchestration requires shared data semantics, identity, policy and escalation.

5. AI economics move from seat licenses to operating leverage

The executive business case should measure total workflow economics: labor capacity released, cycle time, error and rework reduction, avoided delay, quality improvement, incremental throughput and risk reduction—offset by model, integration, validation, governance and change-management costs.

10. Governance, Risk and Decision Rights

10. Governance, Risk and Decision Rights

Agentic systems enlarge the enterprise risk surface because they can do more than generate information—they can take action. With conventional generative AI, the primary risk is often an incorrect or misleading answer that a person can review before acting. An enterprise agent may instead be authorized to retrieve data, call an API or software tool, modify a project or quality record, send a communication, initiate a workflow, or trigger another system or agent. The consequence is that a flawed assumption, incomplete evidence, compromised instruction, or excessive permission can be converted directly into an incorrect action and then propagated downstream. For example, an R&D agent that incorrectly concludes a verification milestone has been met could update project status, trigger manufacturing-transfer work, or feed the error into a Stage-Gate package. Security researchers and standards bodies therefore increasingly emphasize risks specific to agentic systems: excessive agency (more authority than the task requires), identity and permission failures, tool misuse, memory manipulation, data exposure, and cascading multi-agent behavior.[9][10][11] The governance principle is straightforward: agent autonomy should never exceed the organization's ability to detect, contain, reverse, and learn from an incorrect action.

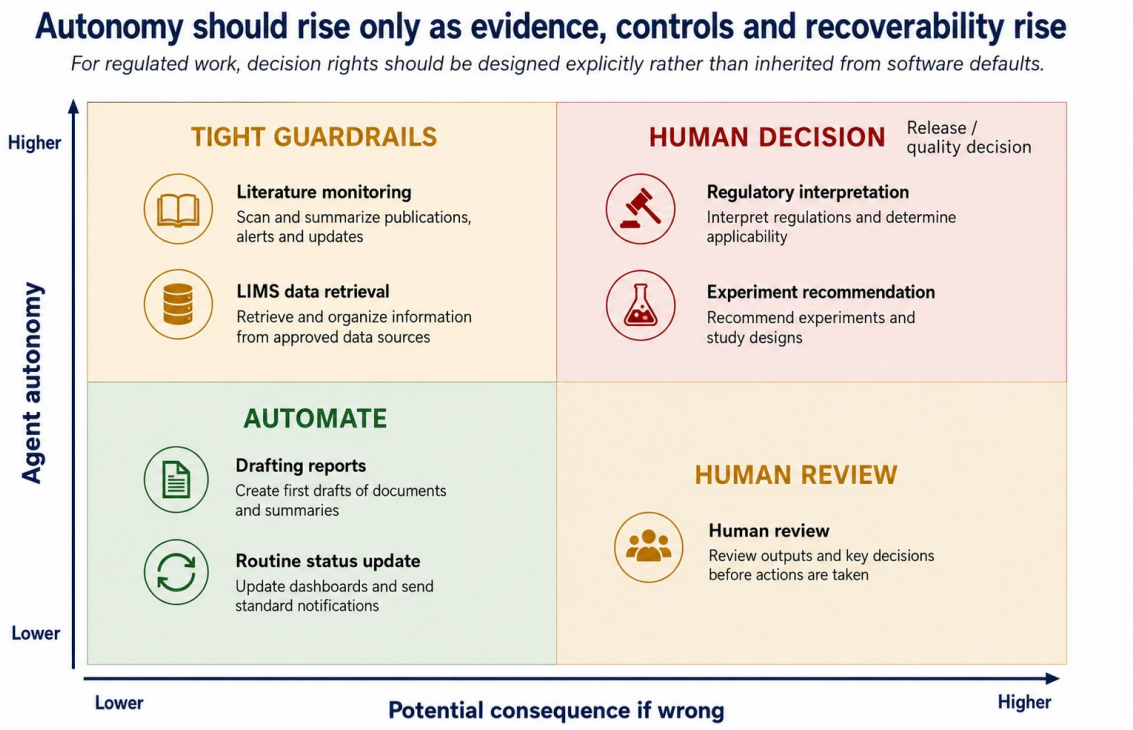

Exhibit 6. Autonomy should be calibrated to consequence and recoverability. Source: CatalystDx framework.

A practical governance hierarchy

Control	Executive question
Purpose & owner	What business outcome is the agent accountable for, and who owns that outcome?
Authority	What may it read, write, spend, communicate, approve or initiate?
Data	Which sources are authoritative, and what privacy or confidentiality restrictions apply?
Evaluation	How is quality tested before and during production? What is the acceptable failure rate?
Human checkpoints	Which actions require review, dual approval or segregation of duties?
Monitoring	Can the organization see tool calls, decisions, errors, overrides and performance trends?
Incident response	How is an agent stopped, rolled back, investigated and restored?
Change control	How are model, prompt, tool, policy and workflow changes assessed and released?

NIST's AI Risk Management Framework emphasizes governing, mapping, measuring and managing AI risk across the lifecycle.[9] ISO/IEC 42001 provides a management-system approach to responsible AI governance and continuous improvement.[10] For organizations deploying agents, those principles need to be extended to operational authority: identity, tool access, logging, permissions, escalation and recoverability.

Regulated life-sciences context

For medical-device and diagnostics organizations, the governance model must also distinguish between AI used to operate the enterprise and AI that is part of a regulated product or that produces evidence supporting regulated decisions. FDA's AI-enabled device guidance emphasizes total-product-lifecycle risk management, documentation and evidence appropriate to safety and effectiveness.[8] The specific regulatory treatment of an enterprise agent will depend on its intended use and context; however, the management discipline is directionally consistent: define intended use, document controls, validate performance for context, manage changes and preserve traceability.

CORE RULE
An agent's authority should never exceed the organization's ability to detect an error, understand what happened, stop propagation and recover safely.

11. A 90-Day Executive Action Agenda

Organizations do not need an enterprise-wide "agent strategy" before learning. They need a disciplined portfolio of bounded experiments that reveal which workflows create value and which operating capabilities are missing.

Weeks	Executive actions	Deliverable
1–2	Select one executive sponsor and establish a small cross-functional agentic-AI working team. Define non-negotiable security, privacy, regulatory and data boundaries.	Governance charter and decision-rights principles.
2–4	Map 10–20 candidate workflows. Score recurrence, value, measurability, data availability, reversibility and consequence.	Prioritized agent opportunity portfolio.
4–6	Choose 2–3 bounded pilots. Define baseline cycle time, labor, quality, failure modes and success metrics before building.	Pilot charters with measurable economics.
6–10	Build with narrow permissions, explicit evaluations, logs and human escalation. Use real workflow owners, not only AI specialists.	Working pilot and evaluation evidence.
10–12	Compare outcome quality and economics with baseline. Document exceptions, security issues, data gaps and adoption barriers.	Scale / revise / stop decisions.
By day 90	Create a 12-month roadmap for workflow orchestration, data, identity, governance, skills and change management.	Agentic transformation roadmap.

Recommended first pilots in a diagnostics company

4. Competitive and scientific intelligence monitoring with traceable sources.
5. R&D program-status and stage-gate evidence assembly using read-only system access.
6. Requirements-to-evidence traceability and missing-evidence detection.
7. Regulatory-guidance monitoring with human-reviewed impact assessment.
8. Quality-system triage or investigation drafting where final decisions remain human.

These use cases are valuable because they are information-intensive and measurable, yet they can begin with limited or read-only permissions. They teach the organization how to build retrieval, evaluation, identity, auditability and escalation before delegating irreversible actions.

12. Agentic Enterprise Assessment Checklist

Use the following checklist in an executive team or functional leadership workshop. A “no” answer is not necessarily a reason to stop; it identifies a capability that must be built before autonomy can scale safely.

Domain	Assessment statement	Yes / No	Priority action
Strategy	We can name the 3–5 business workflows where agentic AI could create the greatest measurable value.	☐ Yes ☐ No	
Strategy	Each candidate workflow has a named executive or process owner.	☐ Yes ☐ No	
Workflow	We have documented the actual end-to-end workflow, including exceptions and informal workarounds.	☐ Yes ☐ No	
Workflow	We can define a measurable completion condition and quality standard.	☐ Yes ☐ No	
Data	The authoritative data and documents needed by the workflow are identifiable and accessible.	☐ Yes ☐ No	
Data	We understand which information may not be exposed to particular models, tools or users.	☐ Yes ☐ No	
Architecture	Agents can be given least-privilege identities and bounded tool permissions.	☐ Yes ☐ No	
Architecture	We can separate read, draft, recommend, write, execute and approve permissions.	☐ Yes ☐ No	

Domain	Assessment statement	Yes / No	Priority action
Evaluation	We have test cases that represent normal, edge-case and adversarial conditions.	☐ Yes ☐ No	
Evaluation	We can quantify agent quality, error rates, escalation rates and human override.	☐ Yes ☐ No	
Governance	High-consequence decisions and regulated approvals have explicit human owners.	☐ Yes ☐ No	
Governance	We can reconstruct what the agent saw, decided, called and changed.	☐ Yes ☐ No	
Security	We have controls for prompt injection, malicious content, tool misuse, credential exposure and excessive agency.	☐ Yes ☐ No	
Operations	We can stop or disable an agent quickly and recover from an incorrect action.	☐ Yes ☐ No	
Change	Changes to prompts, models, tools, policies and workflows have an appropriate release process.	☐ Yes ☐ No	
Economics	We measured the baseline workflow cost and cycle time before deploying the agent.	☐ Yes ☐ No	
Economics	We can calculate value using outcomes—not only hours “saved.”	☐ Yes ☐ No	
People	Employees understand when they remain accountable even when an agent performs the work.	☐ Yes ☐ No	
People	Managers are prepared to redesign roles and measures rather than simply add another AI tool.	☐ Yes ☐ No	
Scale	We have a plan for shared identity, data, observability and governance rather than isolated functional bots.	☐ Yes ☐ No	

SCORING GUIDANCE

0–7 “Yes”: foundation stage. 8–14: pilot-ready with targeted gaps. 15–17: workflow-scale candidate. 18–20: strong foundation for orchestrated multi-agent deployment. The score is directional; a single severe security, regulatory or decision-rights gap can outweigh the total.

13. Conclusion: Rebuild the Workflow, Not Just the Interface

The first wave of generative AI taught organizations that knowledge work could be accelerated through a natural-language interface. The next wave will force a more difficult question: how much of a workflow can software responsibly own?

The answer will vary by process. Some agents will remain narrow and assistive. Others will monitor continuously, coordinate specialized capabilities and execute routine work with little intervention. In regulated and high-consequence environments, human authority will remain essential—but the evidence preparation, coordination and monitoring surrounding those decisions can change dramatically.

The strategic advantage therefore is unlikely to come from possessing a particular model. Models will improve and vendors will change. Durable advantage will come from organizational assets that are harder to copy: clean evidence and data, well-defined workflows, explicit decision rights, evaluation systems, trust, domain knowledge and the ability to redesign work quickly.

For life-sciences leaders, the opportunity is especially significant. R&D organizations have long struggled with fragmented evidence, sequential handoffs, late discovery of readiness gaps and expensive cross-functional coordination. Agentic AI can help create a more continuous operating model in which evidence is gathered as work occurs, risks are surfaced earlier and human decision-makers receive a more complete, current and traceable basis for action.

FINAL EXECUTIVE TAKEAWAY

Do not begin with autonomy. Begin with accountability. Define the workflow, the owner, the evidence, the economics and the decision rights—then introduce the minimum autonomy required to improve the outcome.

14. Preview — Paper No. 2: From Copilots to Agents

Paper No. 2 will move from the organizational level to the design of the agent itself. It will give executives and transformation leaders a practical way to distinguish a chatbot, copilot, task agent, workflow agent and multi-agent system—and to decide which architecture fits a particular business problem.

- The agent design canvas: goal, role, context, memory, tools, policies and success criteria.
- When a single agent is enough—and when multi-agent architectures add unnecessary complexity.
- Tool use, structured actions and enterprise system integration.

- Memory: what agents should remember, what they should forget and why data governance matters.
- Planning and orchestration patterns for long-running workflows.
- Evaluation: how to test an agent before it receives meaningful authority.
- Human-in-the-loop patterns: approve, review, supervise, exception-handle and audit.
- Five diagnostics examples developed from copilot concept to deployable bounded agent.

Working title: “From Copilots to Agents: Designing AI Systems That Can Reliably Perform Enterprise Work.”

References and Further Reading

[1] **AI Agents are Only the Beginning: Deloitte Survey Examines the AI Readiness Gap and Reveals How Enterprises Can Prepare for Agentic Success.** Deloitte, August 12, 2026. https://www.deloitte.com/us/en/about/press-room/deloitte-survey-examines-ai-readiness-agentic-ai-success.html

[2] **2026 Work Trend Index: Agents, Human Agency, and the Opportunity for Every Organization.** Microsoft, May 5, 2026. https://www.microsoft.com/en-us/worklab/work-trend-index/agents-human-agency-and-the-opportunity-for-every-organization

[3] **Capitalize on the AI Agent Opportunity.** Gartner, 2025. https://www.gartner.com/en/articles/2025/ai-agents

[4] **Emerging Tech: Agentic AI Maturity Roadmap.** Gartner, August 8, 2025. https://www.gartner.com/en/documents/6825634

[5] **Towards agentic science for advancing scientific discovery.** Nature Machine Intelligence 7, 1373–1375, September 2025. https://www.nature.com/articles/s42256-025-01110-x

[6] **The Virtual Lab of AI agents designs new SARS-CoV-2 nanobodies.** Nature 646, 716–723, July 2025. https://www.nature.com/articles/s41586-025-09442-9

[7] **Can Agentic AI Improve Workflows and Margins in MedTech?.** Deloitte, November 20, 2025. https://www.deloitte.com/us/en/industries/life-sciences-health-care/blogs/health-care/can-agentic-ai-improve-workflows-and-margins-in-medtech.html

[8] **Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations.** U.S. Food and Drug Administration, January 2025; see FDA digital-health guidance list for subsequent updates. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing

[9] **Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1).** National Institute of Standards and Technology, July 2024; updated

April 2026. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

[10] ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system. International Organization for Standardization. https://www.iso.org/standard/42001

[11] OWASP Top 10 for Agentic Applications. OWASP GenAI Security Project, December 2025. https://genai.owasp.org/2025/12/09/owasp-genai-security-project-releases-top-10-risks-and-mitigations-for-agentic-ai-security/

[12] 2025 Life Sciences Outlook. Deloitte Insights, 2025. https://www.deloitte.com/us/en/insights/industry/health-care/life-sciences-and-health-care-industry-outlooks/2025-life-sciences-executive-outlook.html

About CatalystDx Executive Education

CatalystDx Executive Education is designed to help leaders in diagnostics, medical technology and life sciences translate emerging technology into practical operating models. Future papers in this series will address agent design, agentic R&D, AI-augmented Stage-Gate, multi-agent enterprise architecture, economics and ROI, governance, leadership and implementation roadmaps.

Educational notice. This publication is for executive education and strategic discussion. It is not legal, regulatory, clinical, quality-system, cybersecurity or investment advice. Organizations should evaluate applicable requirements for their specific products, intended uses, jurisdictions, data and operating environments.