msg industry advisors

Headless **Pharma**

Governing agentic AI when the screen is no longer the control layer.

For pharma leaders responsible for AI governance, validation, IT architecture, digital transformation and regulated operations.

Dr. Dennis Janning, Head of Artificial Intelligence, August 2026

CONTENTS

The first part develops the argument. The diagnostic then turns the argument into an assessment tool. The closing pages translate the diagnosis into action and list source references.

Part I: The argument

Part II: The diagnostic

Part III: Decision

I.

The argument

From AI adoption **to headless control**

1. EXECUTIVE SUMMARY

Pharma is not short of AI. It is short of architecture.

Pharma companies are adding AI through vendor modules, copilots, and agentic layers. Each decision can be locally rational. Taken together, they create architectural debt when data access, permissions, memory, execution logic, and auditability stay vendor-specific.

Core Points	Implication
AI adoption is now the floor.	When broad adoption becomes table stakes, the differentiator moves from access to AI toward architecture, governance, and scaled control.
Agents need governed access.	Agentic systems do not depend on screens the way human users do. They act through data, workflow, business logic, policy, and permissions.
The screen can no longer carry attribution alone.	When work moves through APIs, skills, and agent-mediated channels, attribution has to be anchored in the governed data and policy layer.
Cost and validation share one root problem.	Undefined processes and uncontrolled data make AI more expensive to run and harder to defend in audits.
The owned layer is narrow, but essential.	The task is not to rebuild ERP, MES, Veeva, or LIMS. It is to govern data access, policy, attribution, and skills above them.

2. READER LOGIC

From adoption to agentic access to headless control

This whitepaper follows one argument: AI adoption is easier than governing the system that emerges one vendor decision at a time.

1. **Adoption**
   Vendors embed copilots and agents into enterprise applications. Pharma companies buy them because each module solves a local problem.

2. **Fragmentation**
   Those modules bring separate models, memories, permission structures, audit trails, and assumptions about user responsibility.

3. **Agentic access**
   Agents increasingly act on governed data, workflow, and executable business logic rather than helping a human operate a screen.

4. **Headless control**
   The decisive layer becomes the data and policy layer above the systems of record, where permissions, evidence, attribution, and rollback must be designed.

5. **Diagnostic**
   The readiness question is whether the stack can be governed as a set before the next vendor renewal adds another layer.

The next switching event in pharma is a fight over the data and control layer above existing systems. Not simply another cloud or application migration.

3. ADOPTION PROBLEM

Eight AI subscriptions are not a strategy

A mid-size pharma IT estate can already contain multiple AI layers: Veeva agents for quality and regulatory, SAP Joule across ERP, Salesforce Agentforce on the commercial side, Microsoft Copilot, IQVIA AI on market data, MES-related AI capabilities, a document-system copilot, and a team piloting ChatGPT Enterprise.

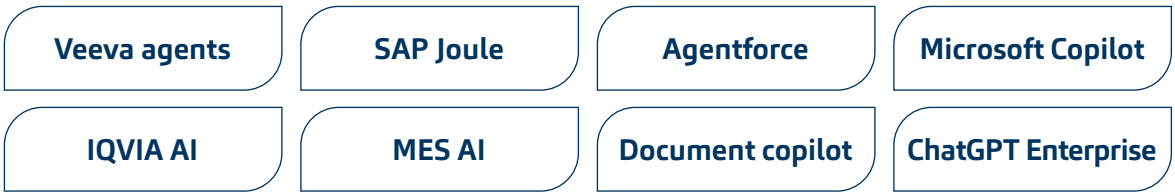

Eight layers · eight permission models · eight memories · eight audit patterns. None of which become an architecture on their own.

None of these decisions has to be wrong in isolation. A capable team can evaluate a capable tool and make a sensible purchase. The problem appears when the estate is viewed as a system: eight AI layers do not automatically resolve into one.

The issue is not vendor AI as such. It is buying AI module by module and leaving the architecture to emerge from the purchase history.

4. ARCHITECTURE

The data and policy layer becomes load-bearing

A copilot helps a person use an application. An agent increasingly wants what sits below the application: governed data, executable business logic, scoped permissions, and the right to act. If that access is mediated through separate vendor APIs and separate audit formats, the agent environment is hard to govern by construction.

This does not mean opening an AI program with a two-year data-cleanup project; that sequencing often creates paralysis. The point is narrower. Before agentic use cases scale, the data and policy layer has to be designed as the load-bearing wall that carries access, semantics, permissions, quality controls, and evidence.

Layer requirement	Why it matters
Data semantics	Meaning should live in the data layer, not only in application code or local process knowledge.
Access patterns	Agents need reliable, governed access paths rather than uncontrolled app-specific shortcuts.
Policy enforcement	Rules must be executable where the agent acts, not only described in a SOP or embedded in a screen workflow.
Evidence and attribution	The record must reconstruct what happened, against which data, under which authority, and with which approved intent.

5. DATA READINESS

FAIR stops being a research-data slogan

For agentic AI, FAIR data becomes an operating condition. Data has to be findable, accessible, interoperable, and reusable not only for human teams, but also for machines that need to locate and use it within defined permissions.

Master data management moves from IT hygiene to quality control. If data is findable but not permissioned, accessible but not attributable, interoperable but semantically inconsistent, the agent does not become autonomous in a useful sense. It becomes expensive middleware with a language model attached.

2.60 → **1.55**

Data quality becomes a safety control

A clinical-summarization study on MIMIC-IV reports fine-tuning on cleaned training data reducing Llama 2 hallucinations from 2.60 to 1.55 per summary. One study does not establish a general law, but the direction is clear: in regulated processes, data quality can become a patient-safety control, not only a data-team metric.

The same constraints that determine whether an agent can be trusted also determine how expensive it is to run. Poorly structured data, unclear semantics, and undefined process steps force the model to compensate for what the operating environment has not resolved. That is where cost control enters the architecture discussion.

Agent readiness requires usable data, controlled access, semantic consistency, and attribution. Missing any one turns autonomy into unmanaged integration.

6. FINANCIAL CONTROL

Token economics exposes process debt

The cost issue is not only model price. It is workload shape. Work on AI token economics points to a pattern CIOs and CFOs need to manage: unit prices may fall, while total AI spend rises as usage, model complexity, and workload intensity increase.

The move to usage-based copilot billing is a procurement and budgeting signal. Input, output, and cached tokens become part of the cost structure. At enterprise scale, model choice, process definition, and data quality become financial controls.

INDIVIDUAL

Vibe to the outcome
Iterate until it feels right.

Strongest model, every time
Reach for the most capable model.

Cost feels free
Marginal cost is invisible.

Individual luxury. Invisible cost, invisible impact.

ORGANIZATION

Robust process and usable data
Defined steps. Controlled data.
Accessible and valuable.

A smaller model is enough
Right model for the job.
Efficient and effective.

Repeatable and auditable
Reproducible output.
Defensible in an audit.

Organizational advantage.
Visible value, sustainable at scale.

The real cost driver is not the model. It is starting from zero, every time.

- **Process reduces tokens**
- **Data reduces waste**
- **Discipline reduces spend**

7. COST AND VALIDATION

Defined steps, controlled data, reconstructable output

At scale, token cost becomes a proxy for process quality. A model on clean input and a defined process can finish the task and stop. A model on uncontrolled input and an undefined process burns tokens compensating for the surrounding system. Early research on agentic workflows also points to high token variance: a workflow that cannot predict its own token burn is hard to budget and difficult to justify.

In a regulated workflow, the same build pays off twice. A process that cannot be reproduced is difficult to validate. An AI-supported deviation summary, validation test draft, or regulatory document workflow needs defined steps, controlled data, reconstructable output, and evidence proportionate to intended use and process risk. This is no longer only an internal design question. Enterprise applications are starting to expose data, workflows, and business logic directly to agents. The screen-based assumptions behind many operating and validation models will increasingly be tested from the outside in. Salesforce made that change visible.

Requirement	Control logic
Cost control requires	Right model for the task; defined input; reusable process; governed data access; predictable workload shape.
Validation requires	Intended use; process risk; controlled data; reconstructable output; evidence trail; review and escalation criteria.
Shared build	Process definition, data quality, permission scope, logging, review criteria, monitoring, and change control.

The cost hedge and the validation case are the same build.

8. SWITCHING EVENT

Salesforce made the screen optional

On April 15, 2026, at TrailblazerDX, Salesforce announced Headless 360. The announcement exposed the CRM platform through APIs, more than 60 MCP tools, and 30 preconfigured coding skills, so that external agents can access data, workflows, and business logic. The architecture signal was explicit.

> **No browser required.**
> The interface is no longer the assumed point of access. An agent connects and acts.

The architecture guidance is the part pharma leaders should read. Headless does not remove governance; it moves governance deeper: into identity, permission design, authorization scope, invocation control, traceability, and token lifecycle. In a browser-based model, the user logs in and the trust model is familiar. In a headless model, an external agent connects and acts through scoped permissions. The company has to prove what it was allowed to do and what it actually did.

Treat Salesforce as a visible market signal. SAP Joule is already moving in this direction conceptually; Veeva, IQVIA, and MES vendors face the same architectural pressure.

9. SYSTEM OF DELEGATION

The system stays, the front door moves

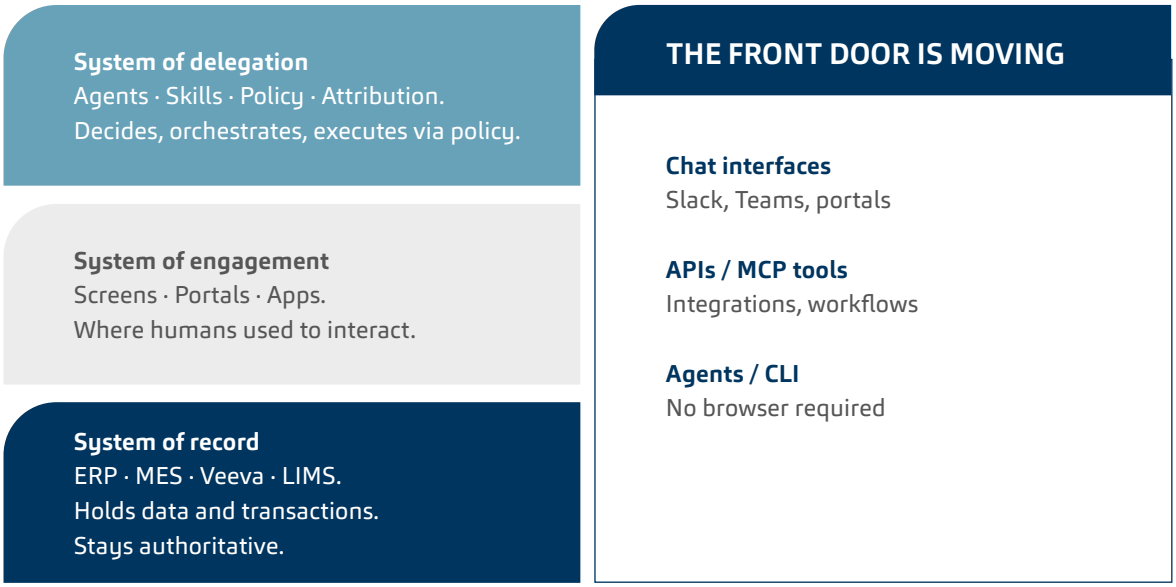

Headless shifts work from engagement to delegation.

The system of record stays authoritative: ERP, MES, Veeva, and LIMS still hold the data, workflows, and transactions. The system of engagement is the screen, portal, app, Teams, Slack, or chat interface where humans used to interact. The system of delegation is the new layer: agents, skills, policy, attribution, and orchestration.

The ERP does not disappear. Nobody is ripping out the system of record. What changes is the interface monopoly: work can be initiated, routed, or rendered through an agent-mediated layer while authoritative records remain in the underlying systems.

10. CONTROL LAYER

When the interface moves, governance must move with it

When the front door moves away from the validated screen, the control layer can no longer be assumed to live in the application interface. The governed data and control layer has to carry authorization, attribution, evidence, and rollback across the systems underneath.

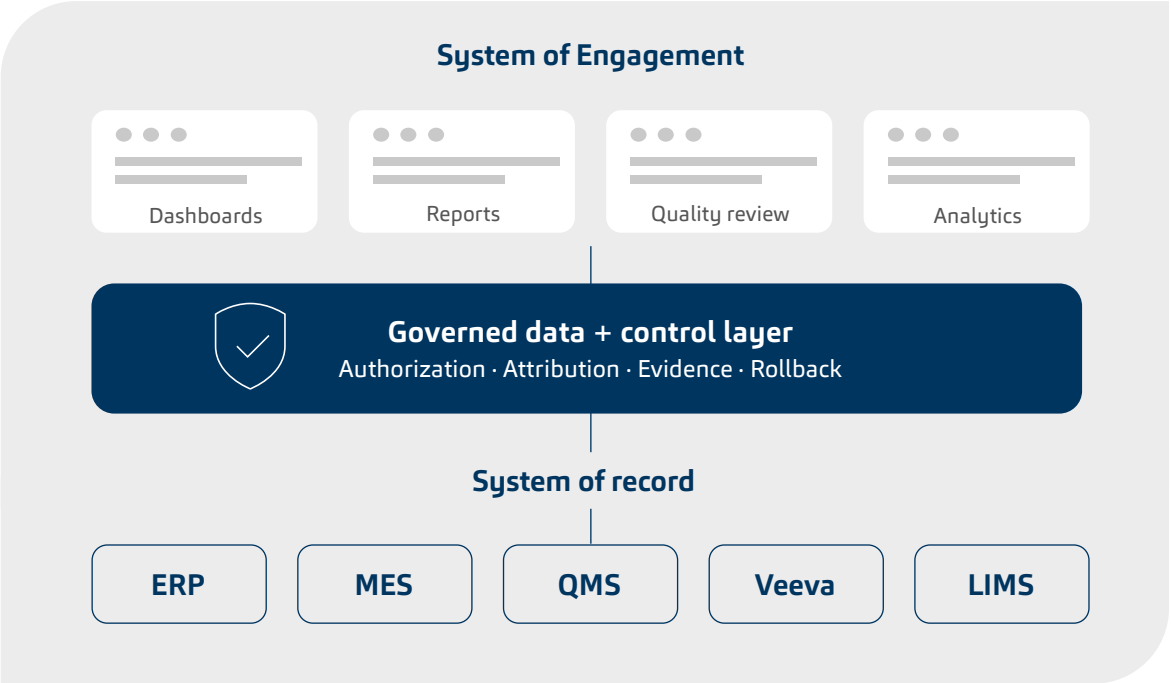

The owned layer sits between the interfaces and the systems of record. Narrow, senior, and load-bearing.

Core question: if an agent acts without opening the validated screen, where does the organization prove authority, intent, evidence, and accountability?

11. VALIDATION IMPACT

Every headless layer is a validation-impact question

Two objections deserve to be taken seriously. The first is the integration argument: pharma CIOs want integrated products that work, not a custom execution layer that becomes a new maintenance burden. The second is the validation argument: in a GxP environment, going headless is not just an architecture change. Changes to schema, permissions, audit structures, or execution logic can trigger change control and validation-impact assessment.

Both objections are valid. Neither supports a wait-and-see position; both support deliberate design. The company is not choosing between change and stasis. Every AI module a vendor adds can affect permissions, data flows, audit events, execution logic, or user responsibilities. The headless shift happens one vendor at a time, whether the enterprise designs for it or not.

Objection	Implication
Integration argument	The clean vendor bundle is attractive until the first audit question asks how the agent acted, under which authority, and with which evidence.
Validation argument	A screen-based, point-in-time validated workflow may be the expected control model. Bypassing it creates impact-assessment work.
Design response	Govern one data and policy layer deliberately, instead of absorbing several uncoordinated vendor layers one after another.

12. PROCUREMENT AND VALIDATION

Enterprise AI procurement can create many validation questions

WHAT A DEPLOYMENT NAMES	WHAT VALIDATION REQUIRES
• Clinical study reports	1. Intended use
• Patient-safety narratives	2. Risk class
• Regulatory submissions	3. System owner
• Drug development as a target area	4. Change control
	5. Audit trail of inputs, prompts, outputs, edits, and approvals

A single enterprise procurement creates a portfolio of regulated workflows, each one assessed on its own.

Illustrative example: Bristol Myers Squibb's rollout of Anthropic's Claude Enterprise to more than 30,000 employees, naming clinical study reports, patient-safety narratives, and regulatory submissions among its target areas (BMS press release and Wall Street Journal, May 2026).

Procurement itself is not the validation trigger. The validation-impact question arises when intended use, process risk, data access, permissions, audit events, execution logic, or user responsibilities change. A broad enterprise deployment can therefore create a portfolio of regulated workflows that have to be assessed individually.

This is why renewal and procurement criteria should include ownership of skills, orchestration portability, audit access to the agent layer, and the ability to operate or evidence the workflow without relying solely on the vendor's black box. FDA and EMA inspect the regulated company, not the software vendor. "The vendor system did it" is not an adequate control answer.

13. OWNED LAYER

What the owned layer actually is

The owned layer is narrower than rebuilding the stack. Pharma companies do not need to rebuild Veeva, ERP, MES, or LIMS. They need a small, senior layer of governance and attribution above the vendors they keep.

Component	Function
Governed data layer	One consistent, quality-controlled substrate where agents access data through consistent semantics and permissions, instead of eight app-specific APIs.
Policy and attribution layer	A thin control plane that decides what an agent may touch, under which authorization, and how action is reconstructed.
Skills as controlled documents	Recurring agent behaviors are written down, versioned, reviewed, and portable across models. Not locked inside a vendor configuration.
Named ownership	Not a project budget handed to one team. A cross-functional charter from CIO or CTO leadership, with Quality, Validation, Data, IT Architecture, and process owners involved.

14. ATTRIBUTION

Attribution has to move below the screen

GxP compliance practice grew up around a person at a screen: who clicked, who approved, who signed. The Annex 11 audit trail can still survive a headless world as a secure, system-generated, time-stamped record. The attribution model layered on top of it is the part that breaks first.

The draft EU GMP Annex 22 is the developing instrument here. As drafted, it focuses on static, deterministic models and indicates that generative AI and LLMs should not be used in critical GMP applications. That position is itself under active reassessment: EMA's 2025 stakeholder consultation showed support for enabling GenAI and LLMs under defined safeguards, and a multistakeholder workshop on 30 June–1 July 2026 gathered further expert input on risk-based controls. Formal guidance is still being shaped while agentic systems already enter enterprise workflows. That makes the company's own control design material.

Not every agent action needs a formal GxP audit trail. But every relevant action needs a clear accountability model with traceability and explainability: what data was used, what changed, which authority allowed it, who approved the outcome, and how the decision can be reconstructed.

If the stack cannot answer which agent acted, under which policy, with which permissions, against which data, then it has a log, not a defensible attribution model.

Attribution Record	Minimum Question
Actor	Which human, system, or agent initiated the action?
Authority	Which policy, role, delegation, or approval allowed it?
Data basis	Which governed data, version, context, or source was used?
Action and evidence	What was created, changed, routed, approved, challenged, or rolled back?

15. REVIEW AND REFRAME

Human-in-the-loop is not a control by itself

Human review becomes a control only when the reviewer has context, competence, time, and a meaningful ability to challenge the output. If the interface makes it easy to click approve, approval becomes routine. Headless makes this harder because the action may happen below or outside the validated screen.

The process therefore needs adversarial review and attribution built into the layer the agent actually touches. Adversarial review means the process is designed to challenge the output, not merely confirm it. The question is procedural before it is technological.

The operating analogy is familiar. A company would not build a sterile fill-finish line by ordering machines from several vendors, each with its own control system and log format, and hope a floor plan emerges from the purchase orders. It would define the architecture first: one layout, one control logic, one place where every action is attributed and every deviation can be caught. The data and policy layer has to play that role for vendor AI.

Review mode	Control quality
Rubber-stamp review	The reviewer sees an output and approves it without enough context, time, or basis for challenge.
Meaningful human review	The reviewer can inspect source data, policy basis, exception logic, and evidence trail.
Adversarial review	The workflow is designed to test the output and surface exceptions before approval.

16. BEFORE THE NEXT AI RENEWAL

Three actions before the next vendor commitment

The next AI renewal is a control decision as much as a commercial one. Each renewal or new module either reinforces a shared data-and-policy layer or adds another isolated permission, memory, audit, and execution pattern. The entry point is not a multi-year architecture program. It is the next decision already on the table.

Today

Inventory AI-enabled modules and pilots. For each, capture model or service, memory, permission model, data-access path, audit trail, and process owner.

This week

Decide where the governed data and policy layer lives, who owns it, and how attribution survives when the screen disappears.

This month

Make execution-layer ownership a contract clause: skill ownership, orchestration portability, audit access to the agent layer, and the ability to operate without the vendor.

Use the diagnostic on the following pages to locate where cost, audit exposure, and vendor dependence are likely to accumulate.

II.

The diagnostic

How headless-ready **is your stack?**

FROM ARGUMENT TO DIAGNOSTIC

Use the diagnostic to locate the control gap

The diagnostic translates the argument into five operational dimensions. It is built for a short working session with IT architecture, Quality, Validation, Data, Procurement, and selected business-process owners. The score is not a label. It is a practical way to identify where the next renewal, pilot, or agentic workflow may increase exposure.

5 dimensions · 10 minutes · score honestly

Score each dimension 0 to 2 and add the points.
The bands on page 22 show what your total means.

0	Not in place	1	Partial / recognized	2	Designed and owned

DIMENSION 1

Architecture before subscription

Before signing the next vendor AI module, does someone map how it touches your data and control layer?

0	Each subscription is judged on its own, in its own business unit.
1	An architecture review exists, but runs after procurement, as cleanup.
2	No AI module is signed without a named owner mapping its data, permission, and audit footprint against the others.

DIMENSION 2

One governed data layer

Do agents reach one governed data layer, or each application's API separately?

0 Every agent reaches into whichever app it sits in. No shared substrate.

1 A data platform exists, but agents still bypass it through app-specific integrations.

2 Agents access a governed data layer with consistent semantics, permissions, and quality controls.

DIMENSION 3

Data, not screens

When you deploy an agent, does it get data access with attribution, or app access through a seat or a screen?

0 Agents operate through human seats and UI automation. Nobody can say what they touched.

1 API access exists, but permissions are coarse and attribution is partial.

2 Scoped data access, least privilege, full action attribution.

DIMENSION 4

Attribution anchoring

Where does Part 11 attribution live: in the data and policy layer, or in screens the agents bypass?

0	Attribution assumes a human at a screen. Agent actions are not separately attributable.
1	The gap is recognized. No design change yet.
2	Attribution, versioning, and rollback live in the data and policy layer, independent of any interface.

DIMENSION 5

Execution-layer ownership

Do you own and control your skills and orchestration, or is it vendor out-of-the-box?

0	Skills and agent behavior are vendor-configured and not portable. You cannot operate without the vendor.
1	Some configuration is documented, most is not.
2	Skills are controlled documents. Your team can develop and run them without the vendor.

YOUR SCORE AND WHAT IT MEANS

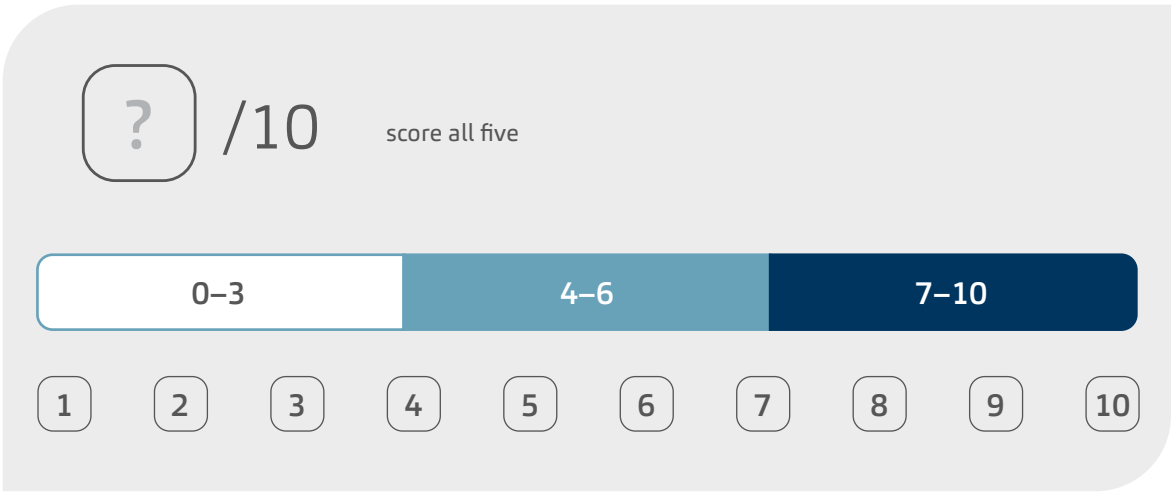

0–3 Adoption, no architecture. The next audit cycle, or the next agent that acts without a clean trail, will expose it. Fix the layer before the next subscription.

4–6 You see the shift; the controls are not in place. One to two quarters of focused work, with a named owner per dimension.

7–10 You treat the data and control layer as the strategy. Pressure-test it each time a vendor in your stack makes its own headless move.

THE PATTERN

Most teams score 2 to 4, not because the people are weak, but because every subscription was sensible on its own, and nobody was given the mandate to govern the set.

III.

Decision

From diagnosis **to the next decision**

CLOSING ARGUMENT

When the screen disappears, someone still owns the answer

The vendors are not the problem. Each vendor is making rational product moves toward agent-ready access, embedded AI, and headless workflows. The risk comes from buying those capabilities faster than the enterprise builds the architecture to govern them. Pharma companies still decide the control model. They can impose one governed data and policy layer on top of vendor AI modules, or they can inherit several control planes with different access paths, memories, permissions, audit assumptions, and validation implications.

The practical test, when an agent acts without opening the screen:

- Which actor or agent acted?
- Under which authority?
- Against which data?
- Through which policy?
- With which evidence?
- Who approved or challenged it?

- **And who owns the answer?**

Headless readiness is the ability to govern action when the interface is no longer the point of control.

DISCUSSION AND NEXT STEPS

Where a focused discussion should start

A useful next discussion starts with the diagnostic score and the next renewal or agentic workflow that could add exposure. The entry point is not a multi-year architecture program. It is the next decision already on the table.

Discussion focus	Typical output
Diagnostic review	Score the stack across the five dimensions and identify where control is already designed, incomplete, or missing.
Control-layer design	Define ownership, data access, permission logic, attribution requirements, evidence trails, and integration principles for agent-mediated workflows.
Validation-impact framing	Translate selected AI scenarios into intended use, process risk, change control, monitoring, and evidence requirements.
Vendor renewal support	Build skill ownership, orchestration portability, audit access, and control-layer criteria into procurement and renewal discussions.

LET'S DISCUSS

Assess and govern headless-readiness before the next AI module becomes part of the validated landscape.

CONTACT:

Dr. Dennis Janning
Head of Artificial Intelligence
dennis.janning@msg-advisors.com

Scan to connect

References

Primary publications behind the argument, for verification and further reading.

1 **Salesforce Headless 360 announcement.** Salesforce, TrailblazerDX, April 15, 2026.

2 **Architecting Trust for Headless Agents.** Salesforce architecture blog.

3 **From AI table stakes to AI advantage: Building competitive moats.** McKinsey (QuantumBlack), May 15, 2026.

4 **GitHub Copilot usage-based billing.** GitHub, effective June 1, 2026.

5 **GMP: AI guidance development, Annex 22.** EMA 2025 consultation closed; position under reassessment; expert workshop 30 June–1 July 2026.

6 **A Data-Centric Approach to Faithful Patient Summaries with LLMs.** Hegselmann et al., CHIL 2024 (PMLR v248); MIMIC-IV.

7 **The FAIR Guiding Principles for scientific data management.** Wilkinson et al., Scientific Data, 2016.

8 **Building the foundations for agentic AI at scale.** McKinsey, April 2026.

9 **How Do AI Agents Spend Your Money? Token consumption in agentic coding tasks.** arXiv, 2026 preprint.

10 **GAMP Guide: Artificial Intelligence.** ISPE.

11 **Open-weight models: data and research.** Epoch AI.

12 **Founders and Forward Deployed Engineers.** AINews / Latent Space, May 30, 2026.

13 **AI token economics for CFOs.** Deloitte, 2026.

14 **Computer Software Assurance for Production and Quality Management System Software.** FDA.

15 **Bristol Myers Squibb rollout of Anthropic Claude Enterprise (30,000+ employees).** BMS press release; Wall Street Journal, May 2026.