ISSUE 03 | NOVEMBER 2026

MAGAZINE FOR COMPLIANCE EXECUTIVES

ComplianceWorxs®

The Record Behind the Decision.

AI at Work in Life Sciences

When employees use AI in consequential work, can the organization show how it affected the decision?

From Assistance to Consequence

Where routine AI use becomes consequential.

"Human Reviewed" Is Not Enough

Why final approval does not prove meaningful review.

What Evidence Leaders Should Expect

A practical framework for AI-supported decisions.

Four Exhibits for Executive Use

Data, decision points, and record requirements.

A Practical Next Step

Test one AI-supported decision.

SPECIAL REPORT

The Business of Compliance

WHAT QUALITY LEADERS ARE REALLY ASKING THE BUSINESS TO FUND

COMPLIANCEWORXS
The Record Behind the Decision.

ISSUE 03 · NOVEMBER 2026

In this issue



FLAGSHIP RESEARCH

AI at Work in Life Sciences

3

When employees use AI in consequential work, can the organization show how it affected the decision?

13 MIN READ



EXECUTIVE PERSPECTIVE

The Business of Compliance: What Quality Leaders Are Really Asking the Business to Fund

11

Quality investment, operational resilience, and the consequences of executive funding decisions.

7 MIN READ



EXECUTIVE STRATEGIC BRIEFING

The Batch Was Released. Can You Defend the Decision?

16

Why approved workflows are not enough — and what it takes to make decisions defensible, retrievable, and inspection-ready.

5 MIN READ

Read and share each article online at complianceworxs.com/magazine/november-2026/



AI at Work in Life Sciences

When employees use AI in consequential work, can the organization show how it affected the decision?

ISSUE 03 · NOVEMBER 2026 · 13 MIN READ · COMPLIANCEWORXS EDITORS

Executive summary

A quality specialist uses an approved AI assistant to summarize evidence for an investigation. A manufacturing team receives an AI-generated warning about a developing process trend. A regulatory writer uses AI to assemble a first draft from prior submissions. A clinical team uses a model to help identify relevant evidence and prepare a protocol.

Each use may be appropriate. Each may save time or reveal something a person might otherwise miss. But if the output materially affects a regulated record or consequential decision, an accountable executive needs answers to a different set of questions:

- What information did the system use?
- What did it produce or recommend?
- What did the employee verify?
- What changed between the AI output and the final decision?
- What uncertainty or risk remained?
- Who authorized the result?

AI is already entering ordinary work across life-sciences companies. The management issue is no longer simply whether employees will use it. It is whether the organization can see where AI enters controlled work and produce a reliable record of how it influenced the outcome.

Read and share online: complianceworxs.com/magazine/november-2026/ai-at-work-in-life-sciences/

The adoption evidence makes that issue urgent. KPMG reported that 48% of life-sciences organizations were strategically integrating AI into core functions and 87% were integrating AI agents into some workflows, products, or services. Deloitte found that only 22% of life-sciences leaders said they had successfully scaled AI and just 9% reported significant returns. ZS found that roughly one-third of employees in each surveyed life-sciences department had approved access to generative-AI tools, but respondents estimated that only about 10% of employees with access used them consistently each week.[1][2][3]

These findings describe an industry in transition. AI presence is high. Access is expanding. Repeated workflow use and measurable value are less mature. The companies moving furthest are not merely buying copilots. They are placing AI inside established workflows, connecting it to approved information, training employees by role, and defining where human review and authorization remain mandatory.[4][5][6]

> *AI presence is high. Access is expanding. Repeated workflow use and measurable value are less mature.*

Yet policies, training, and model controls do not by themselves show what happened in a particular decision. An organization may have an approved platform and still be unable to demonstrate which evidence was considered, how the employee evaluated the output, or why the final action was authorized.

This report is written for executives accountable for quality, regulatory compliance, and other controlled life-sciences functions. It examines what employees are doing with AI, where routine assistance becomes consequential, why "human reviewed" is an insufficient control description, and what evidence leadership should expect when AI contributes materially to a decision.

Its central conclusion is straightforward:

AI can accelerate work without assuming responsibility for it. The organization still needs to preserve the evidence, judgment, risk treatment, and authorization behind consequential AI-supported decisions.

AI is already inside the work

Much of the public discussion about life-sciences AI focuses on drug discovery, model performance, technology partnerships, and long-term transformation. Those developments matter, but they can obscure a more immediate change: employees are beginning to use AI in the work that surrounds everyday decisions.

They use it to find information, summarize documents, compare records, draft content, identify patterns, classify cases, predict outcomes, and recommend actions. The tools may appear as secure chat assistants, specialized copilots, embedded analytics, workflow agents, or features inside systems employees already use.

Large companies provide visible examples. Moderna introduced an internal generative-AI assistant and later expanded access through ChatGPT Enterprise. Sanofi reports that 20,000 employees use its Plai data and decision-support platform every day. Roche developed both a general assistant and domain-specific tools, reportedly reaching approximately 12,000 daily users for RocheChat and more than 5,000 employees for LUMIN.[4][7][8][9]

Companies are also investing in workforce fluency. AstraZeneca made AI learning pathways available in 12 languages and reported more than 12,000 employees participating in tiered certifications by April 2025. Pfizer established an enterprise AI learning program and AI Academy intended to build fluency across employee levels. Sanofi reported broad participation in approved generative-AI learning.[5][6][10][11]

These examples show the direction of travel, not the experience of every company. Many midsize organizations will not build their own enterprise platforms or employ large central AI teams. They will encounter AI through software vendors, Microsoft 365, quality and regulatory systems, clinical platforms, analytics tools, customer systems, and employee experimentation.

That may make the management problem harder to see. AI is not necessarily introduced through one formal enterprise program. It can enter through multiple systems, functions, vendors, and employee practices before leadership has a complete view of where it is being used.

EXHIBIT 1: ACCESS IS NOT ADOPTION

AI adoption should be understood as a progression rather than a yes-or-no condition:

1	**Experimentation**	selected employees test use cases or models.
2	**Approved access**	the company provides a secure assistant or enables AI features in existing software.
3	**Repeated use**	employees rely on AI for recurring tasks rather than occasional exploration.
4	**Workflow integration**	AI is connected to approved data, systems, handoffs, and review requirements.
5	**Scaled value**	the organization can demonstrate better quality, cycle time, capacity, cost, or outcomes without weakening control.

The distinction matters. A license count shows availability. A weekly-user count shows behavior. A completed workflow shows operational adoption. A verified outcome shows value. These measures should not be treated as interchangeable.

The current evidence suggests that access is moving faster than deep integration. KPMG found substantial agent activity alongside limited enterprise maturity in supporting and scaling AI. ZS found a marked difference between approved access and consistent weekly use. Deloitte's scaling and return figures show that broad activity has not yet translated into enterprise results for most organizations.[1][2][3]

For an accountable quality or regulatory executive, the critical question is not where the company sits on an industry maturity curve. It is where AI is already affecting controlled work inside the organization—and whether those uses are visible.

What employees are doing

Although the systems differ, employee uses generally fall into seven task patterns:

- Finding: retrieving documents, prior decisions, scientific evidence, requirements, or internal knowledge.
- Summarizing: condensing studies, records, meetings, cases, trends, or investigations.
- Drafting: producing initial language for reports, protocols, submissions, responses, presentations, or communications.
- Comparing: identifying differences among records, versions, batches, studies, requirements, or outcomes.
- Classifying: sorting cases, deviations, documents, risks, signals, or requests.
- Predicting: estimating properties, outcomes, failures, demand, maintenance needs, or process behavior.
- Recommending or acting: proposing a next step—or, increasingly, executing one within delegated system permissions.

These tasks occur across functions, but their consequences are not equal.

EXHIBIT 2: FROM ROUTINE ASSISTANCE TO CONSEQUENTIAL USE

EMPLOYEE USE	TYPICAL EXAMPLE	CONSEQUENCE INCREASES WHEN THE OUTPUT…
Find	Retrieve relevant procedures or prior records	Determines what evidence is included or excluded
Summarize	Condense an investigation, study, or signal	Becomes the basis for an assessment or approval
Draft	Prepare a report, protocol, submission, or response	Enters a regulated record or external communication
Compare	Identify discrepancies or trends	Drives escalation, disposition, or corrective action
Classify	Categorize a case, deviation, document, or risk	Affects reporting, prioritization, or compliance treatment
Predict	Estimate failure, quality, demand, or outcome	Influences product, patient, trial, or supply decisions
Recommend or act	Propose or execute a next step	Changes a controlled process or commits the organization

A useful governance question is therefore not simply, "Which functions use AI?" It is, "At what point does an AI-assisted task begin to influence a consequential decision?"

Where assistance becomes consequential



Consider a composite example.

A quality team is investigating an unexpected manufacturing result. An approved AI tool retrieves prior deviations, extracts potentially relevant process data, summarizes operator notes, and drafts a preliminary assessment. The investigator reviews the draft, checks selected sources, adds operational context, and recommends that no broader corrective action is required. A quality leader approves the investigation.

The final conclusion may be sound. But an independent reviewer would still need to understand:

Which prior deviations and records the system retrieved.

Whether relevant information was omitted.

How the system linked the evidence to its draft conclusion.

Which sources the investigator actually verified.

What the investigator changed or rejected.

Why the remaining uncertainty was acceptable.

Who had authority to approve the final disposition.

The consequential moment is not necessarily when the employee opens the AI tool. It is when the output begins to shape what evidence is considered, what conclusion is reached, or what action is authorized.

That moment can occur even when AI never makes the final decision. A summary can narrow attention. A classification can determine priority. A recommendation can anchor judgment. A draft can become the structure that reviewers refine rather than reconsider.

This is why a policy statement that "AI does not make decisions" may offer less protection than it appears to. AI can materially influence a decision without holding formal decision authority.

A proportional control model

Not every use needs the same evidence. An employee who uses an approved assistant to improve the wording of an internal meeting invitation does not need a formal decision record. An AI-generated draft supporting a batch disposition, safety assessment, regulatory response, or clinical decision requires substantially more control.

The required evidence should rise with:

The consequence of an incorrect or incomplete output.

The degree to which the employee relies on it.

The difficulty of detecting an error.

The authority delegated to the system.

The reversibility of the resulting action.

The effect on patients, product quality, regulatory evidence, or external claims.

EXHIBIT 3: CONSEQUENCE AND CONTROL

1	**Low-consequence assistance**	
	EXAMPLES	rewriting internal text, summarizing a nonauthoritative meeting, generating presentation ideas.
	EXPECTED CONTROL	approved tool, appropriate data handling, employee verification.
2	**Controlled workflow support**	
	EXAMPLES	searching approved procedures, preparing a report draft, comparing controlled records, identifying potential trends.
	EXPECTED CONTROL	defined use, authoritative sources, trained reviewer, traceable corrections, workflow ownership.
3	**Consequential decision support**	
	EXAMPLES	supporting a quality disposition, safety assessment, clinical conclusion, regulatory response, manufacturing intervention, or external medical claim.
	EXPECTED CONTROL	clear context of use, validated or justified performance, material input and output evidence, meaningful expert review, documented judgment, risk treatment, authorization, and outcome monitoring.
4	**Delegated or agentic action**	
	EXAMPLES	updating controlled systems, initiating communications, changing workflow status, or executing a sequence of actions.
	EXPECTED CONTROL	constrained permissions, identity, approved action boundaries, interruption and escalation mechanisms, complete logging, monitoring, and accountable human ownership.

FDA's 2025 draft guidance proposes a risk-based credibility framework for AI outputs used to support regulatory decision-making, centered on the model's context of use. FDA's later good-AI-practice principles emphasize human-centric design, clear context of use, multidisciplinary expertise, data governance, documentation, performance assessment, and life-cycle management. These principles concern regulatory decision support, but the underlying logic is useful more broadly: evidence and control should match the intended use and risk.[12][13]

Why "human reviewed" is not enough

"Human in the loop" has become a shorthand assurance that AI remains under control. It does not describe what the person actually did.

In the investigation example, the quality specialist may have read the generated assessment and clicked approve. Or the specialist may have traced every material assertion to source data, identified an omitted deviation, corrected a false causal inference, considered an alternative explanation, documented the remaining uncertainty, and changed the proposed action.

Both workflows can be labeled "human reviewed." They are not equivalent.

Meaningful human oversight requires more than presence. The reviewer must:

Have relevant domain competence.

Understand the intended use and material limitations.

Have access to the underlying evidence.

Know which elements require verification.

Have sufficient time to perform a real review.

Be able to change, reject, or escalate the output.

Understand the consequences of acceptance.

Remain accountable for the resulting judgment.

A human click at the end of an opaque automated process does not establish effective oversight.

This matters because fluent AI output can create an illusion of completeness. Reviewers may focus on whether the language sounds reasonable rather than whether the evidence set was complete, the reasoning was appropriate, or contradictory information was excluded.

Organizations should therefore define the human role more precisely. Is the person an editor, source verifier, investigator, decision owner, approver, monitor, or exception handler? What is the person expected to check? What authority does that person retain? What evidence shows that the review occurred?

Training also needs to reflect that role. General AI literacy can help employees use approved tools safely. It cannot replace the function-specific competence needed to review an AI-supported quality, regulatory, safety, clinical, or manufacturing result.

The record leadership should expect

Policies answer what is generally permitted. Validation and model governance answer whether a system is suitable for a defined use. A decision record answers what happened in a particular instance.

For consequential AI-supported work, leadership should expect a record that allows an independent reviewer to follow the material chain from purpose to outcome.

EXHIBIT 4: THE AI-SUPPORTED WORK RECORD

1. PURPOSE
What task or question was the AI used to address? Why was that use appropriate?

2. SOURCE
What authoritative data, documents, records, or context did the system receive or retrieve? What relevant limitations existed?

3. SYSTEM OUTPUT
What material summary, classification, draft, prediction, recommendation, or action did the system produce?

4. REVIEW
Who evaluated the output? What did the reviewer verify, correct, reject, or escalate?

5. JUDGMENT
What reasoning did the responsible person apply beyond the system output? What alternatives or contradictory evidence were considered?

6. RISK
What uncertainty, limitation, exception, or residual exposure remained, and why was it acceptable?

7. AUTHORIZATION
Who had authority to approve the conclusion or action, and when did approval occur?

8. OUTCOME
What happened after the decision? Was performance monitored, and was follow-up required?

This worksheet organizes your notes. It does not make, approve, or recommend a regulated decision. The regulated organization and its authorized personnel remain accountable.

See the Inspection Response Record →

This does not mean retaining every prompt for every use. Nor does it mean surrounding routine work with unnecessary paperwork. It means preserving the material evidence needed to understand an important decision when AI influenced the evidence, analysis, recommendation, or action.

The exact record will vary by context. Some information may already exist in audit trails, source systems, validation packages, controlled records, review histories, or approval workflows. The management task is to determine whether those elements collectively show how the decision was made—or whether the organization would have to reconstruct the story later from system logs and employee memory.

Return to the quality investigation. A complete record would not merely contain the final approved assessment. It would allow a reviewer to identify the sources assembled by the AI, the material draft or recommendation, the employee's substantive review, the reasoning behind the final disposition, the uncertainty accepted, and the authorization.

If those elements are scattered across systems, transient chat sessions, undocumented corrections, and personal recollection, formal approval may exist without meaningful decision visibility.

Conclusion

AI adoption in life sciences is no longer a distant transformation program. It is becoming part of ordinary work.

Employees are using AI to find information, summarize records, draft content, compare evidence, classify cases, predict outcomes, and recommend actions. Companies are providing secure assistants, embedding AI inside established systems, and training employees to use it. At the same time, repeated use, workflow integration, and measurable enterprise value remain uneven.[1][2][3]

For accountable executives, the central risk is not simply that an AI system might be wrong. It is that AI can influence a consequential decision without leaving a record that shows how the influence occurred.

An approved tool is not proof of an appropriate use. A trained employee is not proof of an adequate review. A human approval is not proof of meaningful judgment. A final record is not complete if the organization cannot show the material evidence, output, review, reasoning, risk, and authorization behind it.

The next phase of AI adoption will require more than access and policy. It will require an operating discipline that preserves human accountability as work becomes faster, more automated, and more difficult to reconstruct.

The immediate leadership action is deliberately narrow:

> **Select one consequential AI-supported decision and determine whether its current record shows what the system contributed, what the employee evaluated, what judgment was applied, what risk was accepted, and who authorized the result.**

If the organization cannot answer those questions from the record, it has identified the first control gap to address before AI moves deeper into the work.

FROM COMPLIANCEWORXS

Test one consequential decision.

The Inspection Response Record preserves who authorized it, on what evidence, how risk was weighed, and why it was justified.

See the Inspection Response Record →

NOTES, SOURCES AND CITATIONS

1. **KPMG Global tech report 2026: Life Sciences**
   Open the source ↗
2. **2025 AI trends: Life sciences leaders on data, digital and AI**
   our respondents estimate that only around 10% of employees with access to company-approved gen AI to...
   Open the source ↗
3. **2026 Life sciences outlook | Deloitte Insights**
   Life sciences leaders are generally optimistic for 2026, but see resilience and adaptability as esse...
   Open the source ↗
4. **Digital Transformation and Artificial Intelligence**
   20,000 Sanofians use plai every day—our industry-leading app that delivers real-time, reactive data ...
   Open the source ↗
5. **Upskilling our workforce to thrive in the age of AI**
   Our digital and AI learning pathways are available to all AstraZeneca employees and are available in...
   Open the source ↗
6. **Building an AI-Fluent Organization**
   Resources are available through Pfizer AI Academy, a curated learning platform offering training and...
   Open the source ↗
7. **Our Journey to Becoming a Real-Time AI Organization**
   We developed mChat, Moderna's own generative AI product, modeling tools to automate tasks across dep...
   Open the source ↗
8. **Collaboration with OpenAI: Transforming the Way We Work ...**
   Open the source ↗
9. **How Roche executed a scalable and secure AI strategy**
   Roche built 2 AI apps with OutSystems in 3 months: a daily task companion (12K users) and a domain-s...
   Open the source ↗
10. **2024 Document enregistrement universel**
    Open the source ↗
11. **2025 Impact Report**
    we are investing in our colleagues by offering an enterprise-wide AI learning program to strengthen ...
    Open the source ↗
12. **Considerations for the Use of Artificial Intelligence**
    this guidance provides a risk-based credibility assessment framework that may be used for establishi...
    Open the source ↗
13. **Guiding Principles of Good AI Practice in Drug Development**
    Clear context of use; Multidisciplinary expertise; Data governance and documentation; Model design a...
    Open the source ↗



The Business of Compliance: What Quality Leaders Are Really Asking the Business to Fund

Quality investment, operational resilience, and the consequences of executive funding decisions.

ISSUE 03 · NOVEMBER 2026 · 7 MIN READ · COMPLIANCEWORXS EDITORS

Every company says quality is non-negotiable until a quality decision requires money, time, headcount, validation capacity, or a change in operational priorities. Then the real question appears: does the business understand what it is actually deciding?

For a Quality Assurance or Compliance executive, that is the real commercial problem. The technical and mechanical aspects of compliance are usually well understood inside the Quality function. What is often poorly understood is the business of compliance: how funding, delaying, narrowing, or phasing a quality decision changes the company's operating reliability, release confidence, supply resilience, executive exposure, and ability to withstand scrutiny.

That is why this conversation matters at executive level. When leadership chooses not to fund a quality capability, it is not simply declining a departmental request. It is deciding what level of control, resilience, and defensibility the business is prepared to buy.

Read and share online: complianceworxs.com/magazine/november-2026/the-business-of-compliance/



Compliance is also a business system

It is an operating system that protects the company's ability to manufacture, release, supply, grow, and survive external challenge.

Quality leaders do not need another explanation of CAPA, investigations, GMP training, or deviation handling. They already live inside those mechanics. What is more useful is a clearer way to explain to the rest of the business that compliance is not just a regulatory obligation. It is an operating system that protects the company's ability to manufacture, release, supply, grow, and survive external challenge.

That framing matters because poor quality management does not stay inside the Quality department. FDA's July 2025 white paper describes quality-management investment as a driver of both economic and public-health return, linking stronger quality management with lower costs associated with defects, waste, recalls, and lost revenue, and with greater supply reliability. The same paper argues that these investments can be incremental rather than all-or-nothing, which is important for leaders managing constrained resources across multiple priorities.

That does not mean every proposed quality initiative should be sold as a guaranteed financial win. The stronger claim is narrower and more credible: the business should understand that some compliance decisions protect the conditions under which revenue, supply, and reputation are still possible.

> *"The business should understand that some compliance decisions protect the conditions under which revenue, supply, and reputation are still possible."*

84
INVESTIGATIONS

open and past-due deviation investigations cited at one facility in FDA's amended January 2025 warning letter to Sanofi's Genzyme facility, as of June 2024.

A count, not a percentage. Approximate, as stated in the article.

What the business is actually deciding

When a quality investment request reaches the executive table, it is rarely just about the mechanics written in the memo. A request for investigators, training, oversight capacity, systems, analytics, or stronger process controls usually reflects a deeper business decision about whether the company wants to continue operating with a known weakness.

That is why Quality leaders should frame the discussion in business terms the rest of leadership cannot ignore:

- A deviation backlog is not only an investigation problem; it is a decision about release confidence, management attention, and operational reliability.
- A delayed or weak CAPA is not only a quality-system issue; it is a decision to keep carrying a known failure mechanism.
- Inadequate training is not only a personnel-development gap; it affects execution consistency, deviation recurrence, and inspection defensibility.
- Under-resourcing the quality unit is not simply a headcount choice; it changes the company's ability to identify, explain, and control risk before it affects supply, product disposition, or enforcement exposure.

This is the point many executive teams miss. They believe they are deciding whether to support Quality. In reality, they are deciding what consequences they are willing to own.

> *They believe they are deciding whether to support Quality. In reality, they are deciding what consequences they are willing to own.*

THE FIVE INVESTMENT CATEGORIES

- Investigation capability
- CAPA effectiveness
- Training and knowledge management
- Process control and operational visibility
- Quality-unit authority and capacity

The investments that matter most

This section looks at where real investment decisions tend to sit and why those decisions matter commercially.

The most important investments generally fall into five categories.

1. Investigation capability

If a company cannot investigate well, it cannot explain its own decisions well. FDA's amended January 2025 warning letter to Sanofi's Genzyme facility cited approximately 84 open and past-due deviation investigations as of June 2024 and challenged the adequacy of investigation systems, quality-unit oversight, and the detail behind proposed remediation resources.

For a Quality executive, investment here is not only about adding bodies. It may involve trained investigators, technical subject-matter access, better triage, stronger review discipline, clearer escalation, and sys-

tems that reduce low-value rework. The business outcome is not "more paperwork completed." It is faster, more credible decision-making under pressure.

2. CAPA effectiveness

A CAPA system that closes actions without reducing recurrence gives the business a false sense of control. The 2022 global Quality Benchmarking Study found positive associations between stronger delivery performance and mature practices including CAPA effectiveness, formal risk management, and technical operating discipline.

The practical investment is not merely CAPA software. It includes root-cause capability, effectiveness review, management follow-through, and the discipline to reopen issues that were formally closed but operationally unresolved. That is valuable because repeated failures consume capacity, increase noise in the system, and erode leadership confidence in the quality organization's signals.

3. Training and knowledge management

21 CFR 211.25 requires qualified personnel and continuing CGMP training for people engaged in manufacturing, processing, packing, or holding drug product. ICH Q10 also describes management responsibility for providing the resources needed to implement and maintain the pharmaceutical quality system.

The business case is broader than compliance with training records. Training and knowledge management determine whether the organization can execute consistently, retain judgment as people change roles, and avoid rediscovering the same problem through repeat deviations and slow investigations.

> *Training and knowledge management determine whether the organization can execute consistently, retain judgment as people change roles, and avoid rediscovering the same problem through repeat deviations and slow investigations.*

4. Process control and operational visibility

The benchmarking evidence does not prove causation, but it does show that establishments reporting stronger use of practices such as statistical process control, formal quality risk management, and proactive management of bottleneck equipment tended to show better delivery performance. That matters because quality problems often become business problems only after operational visibility has already failed.

Investment in this category may include monitoring, analytics, real-time process visibility, trend interpretation, spare-part strategy for bottleneck equipment, and the management routines that convert data into action. These are not abstract digital upgrades. They affect the business's ability to detect instability before it becomes waste, delay, deviation volume, or avoidable escalation.

5. Quality-unit authority and capacity

The most overlooked investment is often not technology but decision authority. FDA's Sanofi letter explicitly stated that the quality unit must have the appropriate authority and sufficient resources to carry out its responsibilities and consistently ensure drug quality.

For executive readers, this is central to the business of compliance. A quality unit that lacks authority, time, or credible capacity does not merely become slower. It be-

comes easier for the organization to normalize unresolved risk, postpone difficult decisions, and allow operating pressure to overtake control discipline. That is when commercial urgency starts governing quality decisions indirectly, even if nobody admits it.

What should not be overclaimed

A business case should still preserve mechanical and regulatory reality. That means avoiding claims that are commercially attractive but intellectually weak.

It would be wrong to suggest that every quality investment has a neat or immediate ROI. FDA's 2025 white paper makes a broad economic case for investment, but its model is conceptual and describes company-level economics; it is not a site-specific budget or payback calculator.

It would also be wrong to suggest that framework adoption alone solves the problem. A 2021 analysis of pharmaceutical quality system implementation found improvement in some measures after ICH Q10, but not statistically significant improvement in the effective-management-system category, underscoring that formal adoption and real operating maturity are not the same thing.

That nuance matters. This is not a case for "investing in quality culture" in the abstract; it is an argument for funding real capabilities that change how the business operates.

How Quality leaders should present the decision

The strongest quality-investment cases do not ask the business to admire the importance of compliance. They ask the business to make a clear choice.

A Quality Assurance executive should be able to show:

- What condition exists now.
- What capability is insufficient.
- What investment is required.
- What interim exposure remains if the work is delayed or phased.
- What business consequence becomes more likely if nothing changes.
- How the organization will know whether the investment worked.

That is the shift from technical explanation to executive decision support. It allows Quality leaders to stop arguing that compliance matters and start demonstrating what the company is choosing to protect—or choosing to leave exposed.

The real executive question

The business of compliance begins where the mechanics of compliance stop being enough.

A company can have SOPs, CAPAs, investigations, training records, and dashboards and still misunderstand the business meaning of its own quality decisions. The real question for leadership is not whether Quality has presented a convincing technical case. It is whether the company understands what it is buying when it funds a quality capability—and what it is buying instead when it does not.

That is why the commercial argument matters. Quality leaders do not need to relearn the mechanics; they need to sharpen the argument around them. The leadership team does not need another reminder that compliance is important. It needs help seeing that compliance decisions are business decisions, whether the company treats them that way or not.



The Batch Was Released. Can You Defend the Decision?

Why approved workflows are not enough — and what it takes to make decisions defensible, retrievable, and inspection-ready.

ISSUE 03 · NOVEMBER 2026 · 5 MIN READ · COMPLIANCEWORXS EDITORS

EXECUTIVE READING LENS

An approved batch answers who signed and when. This article examines the more consequential question: can leadership show why the decision was justified from the record that existed at the time?

Executive Summary

A released batch is not always a defensible decision.

Many life-science organizations operate a traditional batch-release workflow successfully — until someone asks a question the workflow was not designed to answer quickly: why was this specific deviation accepted, and can we show the reasoning now?

Read and share online: complianceworxs.com/magazine/november-2026/the-batch-was-released-can-you-defend-the-decision/

KEY TAKEAWAYS

1. **Approval is not the same as a defensible decision.** Signatures close a workflow. A decision record explains the rationale, evidence, risk, and authority.
2. **Rationale is often fragmented.** Key information is distributed across systems, people, and artifacts, making retrieval slow and uncertain.
3. **Start with one decision.** Test one recently completed batch-release decision before changing any system.

PRACTICAL NEXT STEP

Take one recently completed batch-release decision and test whether its record answers, without reconstruction: who authorized it, on what evidence, how risk was weighed, and why it was justified.

The operating reality

Our organization operates through a familiar, traditional batch-release model. Production records are reviewed, exceptions are evaluated, quality approvals are obtained, and batches move forward when the available evidence supports release.

This process has clear strengths. It establishes control, assigns review responsibilities, and documents the formal dispositive of a batch.

But the quality of an approval record is not always the same as the quality of the decision record behind it.

> *the quality of an approval record is not always the same as the quality of the decision record behind it.*

When a batch includes a material deviation — particularly one involving sterile operations, critical process data, environmental conditions, or product-impact assessment — the executive question is not limited to:

"Was the batch approved?"

It becomes:

"Can we immediately show the evidence, risk assessment, expert judgment, authority, and rationale that made approval appropriate?"

That is a higher standard than a completed workflow. It requires a decision record that is retrievable and understandable without reconstructing the story after the fact.

The manual-deviation burden

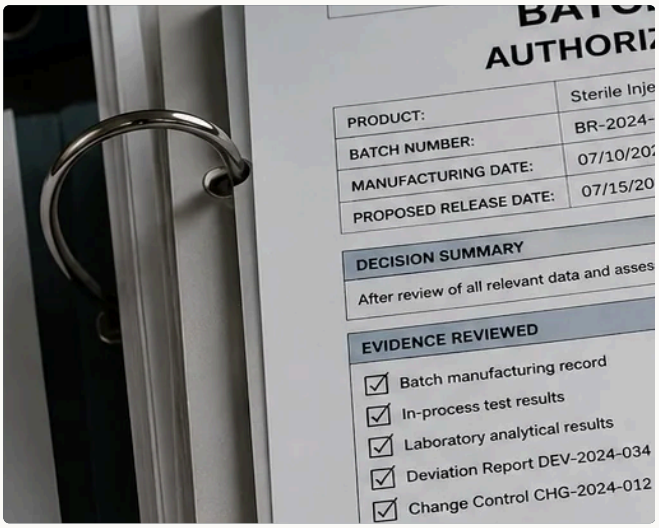

We rely heavily on manual deviation reviews. These reviews often require experienced quality personnel, manufacturing leaders, technical subject-matter ex-

perts, and sometimes senior management to interpret a complex set of facts.

The team may do excellent work. It may assess product risk, review relevant process and analytical information, consider alternatives, determine needed actions, and make a sound release recommendation.

The weakness appears later: the reasoning is often distributed across the process rather than preserved as one coherent decision narrative.

Rationale is often distributed across:

- Batch-record platforms
- Laboratory systems
- Deviation or CAPA tools
- Document repositories
- Email
- Collaboration channels

Each artifact may be valid on its own. Together, however, they can create a difficult-to-navigate evidence trail. An executive or investigator may find the final outcome but still need time — and assistance from several people — to understand why that outcome was justified.

Distributed audit trails create a retrieval problem

Our quality audit trails are distributed across systems. Batch-record platforms, laboratory systems, deviation or CAPA tools, document repositories, email, and collaboration channels may all contain relevant information.

WHY APPROVED WORKFLOW TO DEFENSIBLE DECISION RECORD

The gap is rarely the existence of records. It is the connection among them.



EXECUTIVE TEST

Can a qualified reviewer understand the basis without reconstructing the story?

Figure 1. A defensible decision depends on connecting the record, not merely locating its parts.

This creates a practical difference between having information and being able to use it.

A traditional workflow can demonstrate that:

- The required approvals occurred.
- The workflow moved through the appropriate stages.
- The authorized individuals signed or approved the record.
- Supporting documents existed somewhere within the quality ecosystem.

Why slow reconstruction matters



THE LIMITATION OF A DISTRIBUTED RECORD BECOMES CLEAR WHEN A DECISION IS CHALLENGED — DURING AN INTERNAL REVIEW, MANAGEMENT ESCALATION, OR REGULATORY INSPECTION.

What should be a straightforward explanation can turn into a reconstruction effort. The team must locate documents, revisit email threads, confirm who was consulted, and piece together the rationale across multiple sources.

This takes time, consumes experienced personnel, and introduces uncertainty. It can also create management and reputational exposure if the rationale cannot be demonstrated clearly and consistently.

KEY CONSEQUENCE

A released batch may be compliant. But if the decision cannot be reconstructed quickly, the organization faces unnecessary inspection risk, operational delay, and management exposure.

Why test one decision first

Before making broad system changes, test one recently completed batch-release decision. Pull the record as an inspector would, and see whether the essential elements are present and connected — without anyone reconstructing the story.

WHAT A DEFENSIBLE DECISION RECORD CONNECTS

A challenged decision is easier to defend when the essential elements are present and connected.

One recently completed decision

Shows where the record is complete, connected, and inspection-ready.

A single test can expose gaps such as:

- Missing or inaccessible evidence.
- Unclear risk assessments.
- Inconsistent or undocumented rationale.
- Uncertain decision authority.

One test shows the organization its current exposure and where improvement will matter most — before an inspector asks.

A practical way to test the batch record

Before reconstructing the reasoning behind a release decision, Quality must first be able to establish whether the underlying batch record is complete, traceable, and ready for review.

The ComplianceWorxs CGMP Batch Record Review & Discrepancy Tracker provides a structured way to perform that review. The Excel workbook contains 120 verification points covering the batch record and provides fields to identify supporting evidence, document discrepancies, record reviewer information, and track findings through investigation and resolution.

Rather than relying on an informal review or separate notes, the workbook gives the reviewer one place to identify what was verified, where the supporting evidence resides, what remains unresolved, and which findings require follow-up.

HOW TO USE IT

Use the workbook with one recently completed batch. Review each applicable verification point against the executed record, reference the supporting evidence, and document any discrepancy that requires investiga-

tion or resolution. The resulting review provides a structured picture of the record supporting the release decision.

The workbook does not make the batch-release decision or replace the organization's approved procedures or Quality Unit responsibilities. It helps organize and document the review that precedes that decision.

Where the review identifies a consequential exception, judgment, or accepted residual risk, the Batch Release Authorization Record can then preserve the evidence, judgment, risk evaluation, rationale, and accountable authorization behind the final decision.

Conclusion

A traditional batch-release workflow can produce compliant outcomes. But when the reasoning behind a consequential decision is distributed, it may be difficult to retrieve, connect, and explain under scrutiny.

A defensible decision is not only about having records. It is about having a complete, connected, and retrievable decision record that explains the evidence, judgment, risk, and authorization at the time the decision was made.

Testing one recently completed decision — and preserving it as a single completed record — is a practical way to reduce inspection risk, improve operational resilience, and strengthen executive confidence in high-consequence decisions.

KEY TAKEAWAYS

⚑ **A released batch answers who signed and when.**
A defensible decision record answers why.

Next step

FROM COMPLIANCEWORXS

Batch Release Authorization Record

$75 · One decision · Two practical tools.

Review the record supporting one batch-release decision, then preserve why the decision was justified.

INCLUDED WITH YOUR PURCHASE

$75 — both tools included. Review the batch. Preserve the decision.

complianceworxs.com/batch-release-defense-pack →

NOTES, SOURCES AND CITATIONS

1. **U.S. FDA. Data Integrity and Compliance With Drug CGMP: Questions and Answers. Guidance for Industry. December 2018.**
Open the source ↗

2. **PIC/S. PI 041-1. Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments. 1 July 2021.**
Open the source ↗

3. **U.S. FDA. Quality Systems Approach to Pharmaceutical CGMP Regulations. Guidance for Industry. September 2006.**
Open the source ↗

COMPLIANCEWORXS MAGAZINE

The Record Behind the Decision.

Published monthly for compliance executives.

AI at Work in Life Sciences · Issue 03, November 2026. © ComplianceWorxs. Editorial content is general information and does not constitute legal, regulatory or quality advice for any specific situation.