

# Governing AI in Pharmaceutical Operations

*The AI-Supported Work Authorization Cycle: six gate decisions for whether a defined AI-supported quality workflow may proceed, be narrowed, or be stopped.*

A defensible authorization decision applies to **a defined workflow, not AI in the abstract**. Passing a model test or completing a course does not establish that the combined workflow will produce acceptable work: testing supplies evidence, but the decision also depends on the system's authority, whether people can perform their assigned role, and whether the organization can find and correct failures.

## THE CYCLE: SIX GATES, ONE LOOP

- 0 Inventory & ownership
- 1 Define work & risk
- 2 Set authority & records
- 3 Test the workflow
- 4 Authorize & monitor
- 5 Change, correct, retire

↻ **Restrict, requalify:** findings at Stage 4 (Authorize & monitor) can send scope back to Stage 3 (Test the workflow) for correction.

↻ **Change or failure:** Stage 5 (Change, correct, retire) reassesses scope and authority, looping back to Stage 1 (Define work & risk).

### DECIDE BEFORE THE WORK PROCEEDS

Name the task and accountable owner, examine representative results, and agree on what finding would make the team stop or restrict use. Do this before authorizing a pilot and again before authorizing routine use.

**THE QUESTION EACH GATE HAS TO ANSWER****One question, one decision, per gate****0 Inventory**

*What AI is in use, including unapproved use, and who owns it?*

Register · Contain · Prohibit · Assign owner

**1 Define**

*What task will AI support, and what could failure affect?*

Approve · Narrow · Reject

**2 Authority**

*What may the system read, recommend, change, or approve?*

Authorize bounded functions · Prohibit

**3 Test**

*Can the system and people perform representative work reliably with a reviewer burden they can sustain?*

Authorize pilot · Correct · Narrow · Stop

**Reviewer-burden limit: pass/fail, not a monitoring nicety**

**4 Authorize**

*Does pilot evidence justify routine use, and does performance hold?*

Continue · Expand · Restrict · Investigate · Suspend

**Register tracks the burden limit; threshold fires regardless of omission rate**

**5 Respond**

*What changed, what was affected, and did the response work?*

Reinstate · Narrow · Stop · Retire

■ GATE DETAIL: STAGES 0-1

## 0 Inventory & ownership

### GATE DECISION

Register

Contain

Prohibit

Assign owner

### QUESTION

What AI is in use, including unapproved use, and who owns it?

### EVIDENCE

Use inventory; systems, data, users; named owners.

### OWNER

AI governance lead + quality unit.

Toolkit record: **AI Use Inventory**

## 1 Define work & risk

### GATE DECISION

Approve

Narrow

Reject

### QUESTION

What task will AI support, and what could failure affect?

### EVIDENCE

Intended use, source boundary, measured baseline, risk.

### OWNER

Process owner + quality owner.

Toolkit record: **Gate Decision Record** (light path)

## ■ GATE DETAIL: STAGES 2–3

## 2 Set authority & records

### GATE DECISION

☒ Authorize bounded functions☐ Prohibit

### QUESTION

What may the system read, recommend, change, or approve?

### EVIDENCE

Permissions, approval design, record/signature rules.

### OWNER

System owner, quality owner, IT, data owner.

**Enforced, not just instructed:** the agent's own credentials must technically enforce the envelope, and each approval stays tied to the record state it reviewed. If that state changes before the agent acts, the system determines whether the approval still applies. Another review is required only where the change undermines its original basis.

Toolkit record: **Gate Decision Record:** Authority Envelope

## 3 Test the complete workflow

### GATE DECISION

☒ Authorize pilot☐ Correct☐ Narrow☐ Stop

### QUESTION

Can the system and people perform representative work reliably?

### EVIDENCE

Challenge cases, omissions, false alarms, recovery.

### OWNER

Evaluation owner + quality unit.

**A reviewer's presence is not proof of review:** test whether the review design detects seeded and naturally occurring errors under realistic workload. Automation bias is treated as a control failure mode, not assumed away.

**Reviewer-burden limit:** a pre-set pass/fail condition, not a monitoring nicety. Set it alongside the omission criterion and carry it into Stage 4 and the Authorization Register.

Toolkit record: **Workflow Evaluation Scoring Sheet** + Adjudication Protocol

## ■ GATE DETAIL: STAGES 4–5

## 4 Authorize routine use &amp; monitor

## GATE DECISION

Continue

Expand

Restrict

Investigate

Suspend

## QUESTION

Does pilot evidence justify routine use, and does performance hold?

## EVIDENCE

Pilot results, incidents, overrides, drift, feedback.

## OWNER

Process owner + quality unit.

**Reviewer burden is monitored against the Stage 3 limit here:** the register carries the limit and the current burden against it; the action threshold fires whether or not the omission rate still looks acceptable.

Toolkit record: **AI Authorization Register** + Independent Effectiveness Review

## 5 Change, correct, reauthorize, retire

## GATE DECISION

Reinstate

Narrow

Stop

Retire

## QUESTION

What changed, what was affected, and did the response work?

## EVIDENCE

Change impact, regression, requalification, CAPA.

## OWNER

Change owner, system owner, quality unit.

Toolkit record: **Gate Decision Record** + Management actions

**WHAT “TESTED” HAS TO MEAN****DEVIATION CHRONOLOGY EXAMPLE****59**independent, representative  
holdout cases**0**critical omissions allowed across  
the set**about 4.95%**95% one-sided upper confidence  
bound on the true critical-omission  
rate

These numbers work out the acceptance criterion for one example, a bounded pilot testing an AI-supported deviation chronology. They are not a fixed pass mark for every workflow. Each workflow sets its own criterion sized to its own risks and failure modes: zero critical omissions in 59 independent holdout cases yields a 95% one-sided upper confidence bound of about 4.95%; if one critical omission occurs, at least 93 total cases with no further failures are required to support the same below-5% rate claim; demonstrating a rate below 1% requires 299 zero-failure cases. Cases must be independent, representative of the authorized scope, and held out from any tuning or rehearsal.

**These are separate statistical bounds, not alternate ways to pass:** one omission may still support the same rate claim with a larger sample, but it fails this example's preset zero-omission acceptance criterion.

**Reviewer burden belongs in the same pre-set acceptance criterion as its own pass/fail condition.** The criterion states the review effort per case a reviewer can sustain and still perform the verification authorization depends on. A result above that limit is a failed condition under the pre-set monitoring rule, not an operational inconvenience. The limit and the current burden against it carry into Stage 4 monitoring and the AI Authorization Register.

## THE CONTROL BASIS: REGULATORY REQUIREMENTS AND SUPPORTING GUIDANCE, NOT A REPLACEMENT

21 CFR 211.22(c) Quality-unit responsibility

21 CFR 211.25 Qualified personnel

ICH Q9(R1) Quality risk management

ICH Q10 Pharmaceutical quality system

21 CFR 211.68 / Part 11 Records &amp; signatures

FDA data integrity / PIC/S PI 041-1 Provenance

GAMP Guide: AI / GAMP 5 Lifecycle practice

## WHAT RUNS THE CYCLE: THE AI WORK AUTHORIZATION TOOLKIT (V4.2.1, EIGHT OPERATING ARTIFACTS)

1

### AI Use Inventory

Identifies actual use, ownership, status, and disposition.  
Does not authorize.

2

### Gate Decision Record

Records the decision to authorize a bounded workflow,  
supported by the required evidence and approvals.

3

### Workflow Evaluation Scoring Sheet

Records test evidence, critical-failure handling, holdout  
integrity, and gate recommendation.

4

### Adjudication Protocol

Freezes the critical-omission classification method,  
calibrated and blinded, before the holdout is scored.

5

### AI Authorization Register

Aggregates current state, Decision history, and  
Management actions. Indexes; does not replace GDRs.

6

### Independent Effectiveness Review Record

Tests whether authorization, monitoring, and actions  
actually work in practice.

7

### User Guide

Explains the connected control sequence, artifact  
crosswalk, and implementation boundary.

8

### Role Quick-Start Cards

Gives each of the eight role groups an actionable starting  
point.

No single artifact authorizes a workflow or demonstrates continuing effectiveness. The eight operate as one control system, alongside the companion white paper (v4.2), a Release Notes / Verification Summary, and the v4.2.1 correction notice.

## WHO OPERATES IT: EIGHT ROLE GROUPS, ONE CONNECTED SEQUENCE

### 1. Process owner

*Own the work and acceptable outcome.*

**Do:** Define the bounded task, baseline, downstream decision, and record path.

**Never:** Authorize through an inventory or register entry alone.

### 2. Quality owner / unit

*Decide whether evidence and controls support the regulated use.*

**Do:** Approve thresholds, critical-failure mapping, and the Adjudication Protocol before results are examined.

**Never:** Approve after the fact.

### 3. System / data / security / privacy owners

*Define and enforce what the workflow may access and do.*

**Do:** Maintain and technically enforce the approved Authority Envelope.

**Never:** Rely on a prompt to prohibit an available action.

### 4. Evaluation owner

*Run the scoring sheet and produce Stage 3 evidence.*

**Do:** Freeze the rubric and calibrate on a seeded collection before the protected holdout is scored.

**Never:** Confirm your own mapping as the only reviewer.

### 5. Assessor / operational reviewer

*Judge the evidence independently and document the basis.*

**Do:** Complete the independent assessment before seeing the AI conclusion, when the protocol requires it.

**Never:** Score tone or polish instead of evidence.

### 6. AI governance / portfolio owner

*Keep organizational evidence complete, current, and traceable.*

**Do:** Reconcile the Use Inventory, Authorization Register, and Decision history.

**Never:** Replace local process and quality owners.

### 7. Independent effectiveness reviewer

*Test whether authorization and monitoring controls work in practice.*

**Do:** Test actual use against the Authority Envelope and authorized scope.

**Never:** Review work you designed or operate without disclosed safeguards.

### 8. Management review chair

*Convert organizational evidence into accountable decisions and resources.*

**Do:** Decide whether to restrict, suspend, requalify, resource, or retire affected workflows.

**Never:** Accept a portfolio average that hides a failed subgroup.

## WHAT THIS METHOD DOES NOT ESTABLISH

☐ That every AI-supported workflow is acceptable.

☐ A replacement for validation or existing GMP responsibilities.

☐ Real-world performance of the Cycle or its worked examples.

☐ Authorization without locally approved evidence and an accountable quality decision.