2026 HOSPITAL PHARMACY RXBUYER CONFERENCE SPECIAL EDITION

AI Guardrails for Pharmacy Buyers

A leadership prompt guide for evaluating, governing, and operating AI safely

42 practical prompts | 5 risk categories | 1 approval gate | 90 days to action

Paul J. Swider

CEO & Chief AI Officer, RealActivity | August 2026

Use AI faster—without outsourcing accountability

This guide turns the conference session “AI Risk Management for Pharmacy Leaders: Building Guardrails Before You Need Them” into a working leadership toolkit. It also extends the presentation into vendor evaluation, validation, procurement, monitoring, ROI, agentic AI, and board communication.

> **THE CORE PRINCIPLE** AI may help draft, compare, summarize, test, and plan. Accountable people still decide, approve, verify, document, and intervene.

What you will be able to do

- Identify five categories of AI risk in pharmacy and apply one shared language.
- Run disciplined vendor due diligence across data, transparency, accountability, and exit.
- Move every AI use through one approval gate—with no shadow adoption.
- Build role-based AI literacy without requiring a technical background.
- Prepare incident-response and near-miss learning before the first event.
- Launch a practical 90-day governance operating rhythm.

How to use this guide

1. Choose a prompt based on the decision you need to make—not the tool you want to use.
2. Replace bracketed placeholders and attach only approved, minimum-necessary information.
3. Ask the AI to label unknowns, cite sources, and show assumptions.
4. Use the leadership check before acting. High-risk outputs require qualified human review.
5. Save the final decision, evidence, owner, and date in your normal system of record.

Before you prompt: the pharmacy safety compact

A prompt is part of a controlled workflow—not a waiver of clinical, privacy, security, legal, or professional responsibility.

TIER	EXAMPLES	DEFAULT ACTION
GREEN	Public, synthetic, or fully de-identified material in an approved tool.	Normal review; verify important facts.
AMBER	Internal operational, pricing, employee, workflow, or confidential material.	Use only approved enterprise tools and access controls.
RED	Patient-specific data, credentials, secrets, regulated decisions, or automated clinical action.	Do not proceed without explicit approval, minimum necessary data, and required human controls.

The P.H.A.R.M.A.C.Y. prompt pattern

	ELEMENT	ASK YOURSELF
P	Purpose	What decision, artifact, or outcome is needed?
H	Human role	Who are you asking the AI to emulate, and who remains accountable?
A	Approved assets	Which policies, evidence, data, or files may it use?
R	Rules and risks	What must it avoid, flag, verify, or escalate?
M	Measures	What does good look like, and what thresholds matter?
A	Audience	Who will read or act on the output?
C	Constraints	Scope, time, format, exclusions, privacy, and uncertainty.
Y	Yield	Exact deliverable: table, memo, checklist, script, plan, or questions.

REUSABLE VERIFICATION TAIL "List assumptions and unknowns. Separate facts from inferences. Cite the source for each material claim. Identify what a qualified human must verify before this output is used."

Contents

MODULE 1

Five Categories of AI Risk

Clinical and patient safety • Data and privacy • Model and accuracy • Operational and workflow • Compliance and regulatory

A bad AI output in pharmacy can become a patient-safety event—not merely a bad document.

LEADERSHIP LENS Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

The five-category risk model

CATEGORY	EXAMPLES	CONTROL QUESTIONS
Clinical & patient safety	Dosing, interactions, substitutions, contraindications, prioritization	Pharmacist verification; local validation; stop criteria
Data & privacy	PHI exposure, training leakage, access, retention, subprocessors	Minimum necessary; approved environment; BAA/terms review
Model & accuracy	Hallucination, drift, bias, missingness, out-of-scope use	Thresholds; subgroup tests; version and drift monitoring
Operational & workflow	Automation error, over-reliance, downtime, queue failure, skill decay	Human checkpoints; downtime and recovery; workload metrics
Compliance & regulatory	HIPAA, CMS/ONC, FDA/device status, state board, documentation	Qualified review; intended-use clarity; evidence and audit trail

CONSEQUENCE OVERLAY Reputational harm can follow failure in any category. Track it as an impact, while fixing the underlying clinical, data, model, operational, or compliance cause.

01 AI Use Inventory Interview

USE WHEN you suspect AI is already being used informally | **WORKSPACE** Copilot / ChatGPT / Claude | **RISK** MODERATE

Have ready: *department roster, workflow list, approved-tool policy*

COPY / PASTE PROMPT

Act as a neutral healthcare operations interviewer. Help me discover where AI is already being used in our pharmacy department without creating a punitive tone. Create: (1) a 12-question interview guide for pharmacists, technicians, analysts, buyers, and leaders; (2) a short anonymous survey; (3) an inventory table with use case, tool, data touched, decision affected, user, frequency, approval status, and owner; and (4) language that makes clear that the purpose is safety and support, not discipline. Flag any use involving patient-specific data, medication decisions, external sharing, automated action, or unsupervised output for immediate review. Do not assume a tool is safe because it is already licensed by the health system.

LEADERSHIP CHECK *Run discovery through your privacy/compliance-approved channel; do not ask staff to paste PHI into the survey.*

02 Five-Category Risk Scan

USE WHEN you need a consistent first-pass risk assessment | **WORKSPACE** Word / secure enterprise AI | **RISK HIGH**

Have ready: *use-case description, users, data flow, intended decision*

COPY / PASTE PROMPT

Act as an interdisciplinary AI risk review team for hospital pharmacy. Assess the proposed use case below across five categories: clinical and patient safety; data and privacy; model and accuracy; operational and workflow; compliance and regulatory. For each category, identify hazards, affected people, severity, likelihood, detectability, existing controls, missing evidence, and a recommended risk rating. Add a separate reputational consequence overlay. Finish with: stop / redesign / limited pilot / approve with controls; the three conditions that must be satisfied before proceeding; and the accountable human owner. If evidence is missing, say “unknown” and ask for it rather than inventing an answer. Use case: [PASTE DE-IDENTIFIED DESCRIPTION].

LEADERSHIP CHECK *Treat the output as a draft risk register, not an approval. Clinical, privacy, security, legal, and operational owners must validate it.*

03 Clinical Safety Failure-Mode Workshop

USE WHEN AI could influence medication-related decisions | **WORKSPACE** Teams / Word | **RISK HIGH**

Have ready: *workflow map, intended users, known failure examples*

COPY / PASTE PROMPT

Facilitate a prospective failure-mode workshop for [USE CASE]. Map the medication-use workflow from trigger to final action. At each step, identify how an AI output could be wrong, incomplete, delayed, biased, misapplied, or trusted beyond its intended role. Include dosing, interactions, substitutions, contraindications, formulary status, patient-specific context, alert fatigue, handoffs, and downtime. For each failure mode, propose a prevention control, a detection control, a human verification point, an escalation threshold, and evidence to collect during a pilot. Separate “AI informs,” “AI recommends,” and “AI acts” because the required safeguards differ. Return a concise FMEA-style register and the five highest-priority test scenarios.

LEADERSHIP CHECK *A pharmacist with relevant domain authority must define acceptance criteria and retain final clinical judgment.*

04 PHI and Data-Flow Map

USE WHEN a tool may create, receive, maintain, or transmit ePHI | **WORKSPACE** Visio / Word / secure AI | **RISK HIGH**

Have ready: *architecture sketch, interfaces, vendor data terms, retention details*

COPY / PASTE PROMPT

Act as a healthcare privacy and security analyst. Build a data-flow inventory for [USE CASE] from source to deletion. Identify each system, interface, user role, data element, purpose, storage location, transmission path, retention period, subprocessors, logging destination, and access control. Mark where ePHI, credentials, proprietary pricing, employee data, or confidential contract data may appear. Then list unanswered questions for the vendor and internal teams, including BAA coverage, tenant isolation, encryption, training use, human review by vendor personnel, backup retention, deletion verification, and incident notification. Conclude with a minimum-data version of the workflow and a list of data that should never enter an unapproved AI tool.

LEADERSHIP CHECK *Have privacy, security, and contracting teams verify the actual architecture and legal terms.*

05 Model Validation Plan

USE WHEN a vendor claims accuracy or performance | **WORKSPACE** Excel / Word | **RISK HIGH**

Have ready: *intended use, target population, baseline, vendor metrics*

COPY / PASTE PROMPT

Design a local validation plan for [AI SOLUTION] used for [INTENDED USE]. Define the target population, excluded populations, comparison baseline, sample strategy, gold standard, primary and balancing metrics, acceptable error thresholds, subgroup checks, silent-run period, reviewer qualifications, adjudication process, and stop criteria. Distinguish discrimination from calibration and operational usefulness. Require evaluation of false positives, false negatives, missing-data behavior, out-of-distribution cases, changes over time, and differences between the vendor's validation setting and ours. Provide a results table template and a go/no-go decision rubric. Do not claim that vendor evidence proves local performance.

LEADERSHIP CHECK *Predefine thresholds before seeing results; otherwise the team can rationalize weak performance after the fact.*

06 Workflow Dependency and Downtime Test

USE WHEN automation may become operationally critical | **WORKSPACE** Word / tabletop session | **RISK MODERATE**

Have ready: *current SOP, staffing model, downtime process*

COPY / PASTE PROMPT

Analyze how [AI TOOL] could create hidden operational dependency. Compare the current workflow, the AI-assisted workflow, and a downtime workflow. Identify new failure points, manual workarounds, staffing assumptions, alert queues, integration dependencies, and tasks that may atrophy if staff stop practicing them. Create a downtime playbook with trigger, owner, communication, backlog recovery, reconciliation, and return-to-service checks. Add a quarterly resilience drill and measures for recovery time, backlog, medication delays, overrides, and near-misses. Flag any step where the tool's absence could impair patient care or regulatory documentation.

LEADERSHIP CHECK *Resilience is part of the business case. Price the manual fallback and recovery burden.*

MODULE 2

Vendor Due Diligence

Evidence over promises: data governance, transparency, accountability, implementation, and exit

Ask vendors to show what happens when the model is uncertain, wrong, unavailable, or changed.

> **LEADERSHIP LENS** Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

The three evidence questions

6. What exactly is the intended use—and what is explicitly out of scope?
7. What evidence shows the product works in a setting and population like ours?
8. What happens when the tool is wrong, changes, fails, or leaves?

07 Executive Vendor Question Set

USE WHEN preparing for a first vendor meeting | **WORKSPACE** OneNote / Word | **RISK** MODERATE

Have ready: *use case, strategic objectives, known constraints*

> **COPY / PASTE PROMPT**
> Create a pharmacy-buyer interview guide for an AI vendor. Organize questions under: intended use and exclusions; customer workflow; data governance; model development and validation; human oversight; implementation; integration; monitoring and drift; cybersecurity; accessibility and bias; incident response; accountability; pricing; references; and exit. For each question, state what a strong answer sounds like, what evidence to request, and one red flag. Prioritize the 15 questions that most quickly distinguish operational maturity from marketing. End with five live-demo scenarios that force the vendor to show limitations, error handling, auditability, and user override—not just the happy path.

LEADERSHIP CHECK *Ask for artifacts, not assurances: validation reports, logs, architecture, SLA language, references, and contract terms.*

08 Vendor Evidence Gap Analyzer

USE WHEN reviewing a deck, proposal, or RFI response | **WORKSPACE** Word / secure AI | **RISK** MODERATE

Have ready: *vendor materials with confidential data removed or approved*

COPY / PASTE PROMPT

Review the vendor material below as a skeptical hospital pharmacy buyer. Extract every claim about safety, accuracy, ROI, implementation time, integration, compliance, adoption, and outcomes. For each claim, classify it as demonstrated, supported but incomplete, marketing assertion, or not assessable. Identify the population, setting, comparator, sample size, time period, exclusions, and source when provided. Create an evidence-gap list and follow-up questions. Highlight absolute versus relative improvements, selected case studies, ambiguous denominators, and metrics that do not reflect patient or operational outcomes. Do not infer evidence that is not present. MATERIAL: [PASTE OR ATTACH APPROVED CONTENT].

LEADERSHIP CHECK *Keep the original source beside the analysis so reviewers can trace every conclusion.*

09 Data Governance and BAA Deep Dive

USE WHEN the solution will handle sensitive information | **WORKSPACE** Word / procurement workspace | **RISK HIGH**

Have ready: *vendor privacy addendum, BAA, data-flow diagram*

COPY / PASTE PROMPT

Act as a cross-functional reviewer preparing questions for privacy counsel and contracting. Analyze the vendor's proposed data handling for [SOLUTION]. Create a diligence checklist covering data categories, purpose limitation, minimum necessary use, hosting region, encryption, identity and access, logs, retention, deletion, backups, training and fine-tuning, prompts and outputs, vendor human access, subprocessors, model providers, breach notification, audit rights, and termination. Identify where a BAA or other agreement may be relevant, but do not make a legal conclusion. Produce: unresolved questions, evidence requested, proposed operational controls, and contract topics for counsel to evaluate.

LEADERSHIP CHECK *Only counsel and privacy leadership should decide contractual sufficiency; this prompt is issue-spotting.*

10 Transparency and Limitations Interview

USE WHEN the product influences recommendations or prioritization | **WORKSPACE** Word / vendor meeting | **RISK HIGH**

Have ready: *intended use, vendor technical documentation*

COPY / PASTE PROMPT

Build a structured interview to understand how [MODEL] produces and communicates outputs. Ask about intended users, intended population, decision role, input features, missing data, training and validation populations, external validation, performance by subgroup, confidence or uncertainty, known limitations, out-of-scope use, explainability, update process, versioning, drift detection, human factors, and how users learn that a model changed. Request a model card or equivalent evidence. Add a section asking whether the solution is part of certified health IT, a medical device, neither, or uncertain—and what documentation supports that position. Conclude with a buyer-facing transparency scorecard.

LEADERSHIP CHECK *Regulatory status and certified-health-IT obligations require qualified review; do not accept a sales label as proof.*

11 Accountability, Audit, and Exit Terms

USE WHEN negotiating responsibility before purchase | **WORKSPACE** Word / contract review | **RISK HIGH**

Have ready: *draft agreement, SLA, security schedule*

COPY / PASTE PROMPT

Create a negotiation issue list for [SOLUTION] focused on operational accountability. Cover: named responsibilities; implementation acceptance; performance commitments; service levels; audit logs; model and feature changes; notice and approval rights; incident cooperation; evidence preservation; correction timelines; customer data rights; subcontractors; insurance; indemnity topics for counsel; termination assistance; data export; deletion certification; business continuity; and transition support. For each issue, provide the business rationale, a buyer-preferred position, a fallback position, and the operational owner who must approve compromise. Do not draft legal advice; write clear issue-spotting language for counsel and procurement.

LEADERSHIP CHECK *Make sure the people who must operate the contract review the terms—not only legal and sourcing.*

12 Live Demo Stress Test

USE WHEN planning a vendor demonstration | **WORKSPACE** Teams / scoring worksheet | **RISK** MODERATE

Have ready: *realistic de-identified cases, evaluation criteria*

COPY / PASTE PROMPT

Design a 45-minute scripted demonstration for [SOLUTION] using pharmacy-realistic, de-identified scenarios. Include: normal case; missing data; contradictory data; unusual population; workflow interruption; integration failure; user override; wrong output; audit-log review; downtime; and post-update comparison. For each scenario, specify what the vendor must show, expected user behavior, scoring criteria, and a red flag. Reserve time for the buyer—not the vendor—to choose an unannounced scenario. End with a structured scorecard for clinical safety, usability, transparency, recovery, integration, and evidence quality.

LEADERSHIP CHECK *Never use real patient data in a sales demo unless your organization has explicitly approved the environment and process.*

13 Reference Call Script

USE WHEN checking whether outcomes survived real implementation | **WORKSPACE** OneNote / Word | **RISK** LOW

Have ready: *vendor reference list, your implementation concerns*

COPY / PASTE PROMPT

Create a candid 25-minute customer reference interview for a hospital pharmacy leader using [SOLUTION]. Ask about original problem, selection process, implementation timeline, staffing, integration, training, adoption, overrides, errors, near-misses, downtime, vendor responsiveness, model changes, measurable outcomes, total cost, unexpected work, and whether the customer would buy again. Include follow-up probes that discourage polished yes/no answers. Finish with a comparison template that separates vendor-selected references from independently sourced peers and records differences in organization size, EHR, formulary complexity, and use case.

LEADERSHIP CHECK *A reference proves fit for that organization, not yours. Compare context before generalizing.*

14 RFI Scoring Matrix Builder

USE WHEN turning requirements into comparable evidence | **WORKSPACE** Excel / procurement system | **RISK** MODERATE

Have ready: *requirements, weights, mandatory thresholds*

COPY / PASTE PROMPT

Convert our objectives for [USE CASE] into a weighted AI procurement scorecard. Use categories: patient and clinical safety; workflow fit; integration; data governance and security; model evidence; usability and accessibility; implementation capability; monitoring and support; commercial value; accountability and exit. Define measurable criteria, evidence required, scoring anchors from 0 to 5, weights totaling 100%, mandatory pass/fail gates, and evaluator roles. Include a confidence score so weak evidence cannot earn a strong rating. Add sensitivity analysis instructions to show whether the winner changes when weights shift. Avoid criteria written to favor a known vendor.

LEADERSHIP CHECK *Approve the criteria and weights before opening final proposals.*

MODULE 3

One Approval Gate

Intake • Risk review • Approve and scope • Pilot and monitor • Ongoing review

Every tool passes the same gate. Convenience is not a control.

LEADERSHIP LENS Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

One approval gate

STEP	MINIMUM OUTPUT
1. Intake	One place to log the problem, tool, data, users, action, owner, and benefit.
2. Risk review	Score across five categories and route to the right reviewers.
3. Approve & scope	Define allowed use, prohibited use, controls, owner, and expiration.
4. Pilot & monitor	Test against baseline, thresholds, stop criteria, and balancing measures.
5. Ongoing review	Monitor, manage changes, capture incidents, and reapprove on cadence.

15 One-Page AI Intake

USE WHEN a staff member proposes a new AI use | **WORKSPACE** Forms / SharePoint / Word | **RISK** MODERATE

Have ready: *requester, problem statement, proposed tool*

COPY / PASTE PROMPT

Create a concise AI-use intake form for hospital pharmacy. Capture: problem to solve; current workflow; proposed users; decision or action affected; patient-care relevance; data involved; systems and interfaces; vendor or model; expected benefit; success measure; known alternatives; human oversight; urgency; funding; and accountable sponsor. Add screening questions for PHI, automated action, clinical recommendation, external sharing, model training, patient-facing output, employee decisions, regulated functionality, and shadow deployment. Route responses into low, moderate, or high review intensity and explain the rationale. Keep the form completable in under 10 minutes.

LEADERSHIP CHECK *A short intake increases compliance; reviewers can request detail after the risk signal is visible.*

16 Risk Triage and Review Routing

USE WHEN deciding who must review an intake | **WORKSPACE** Power Automate / Word | **RISK** HIGH

Have ready: *completed intake, committee roles, risk appetite*

COPY / PASTE PROMPT

Using the completed AI intake below, produce a review-routing recommendation. Map each risk signal to the required reviewers: pharmacy operations, clinical pharmacy, informatics, patient safety, privacy, security, legal, compliance, procurement, accessibility, finance, and executive sponsor. Identify whether the request can be rejected quickly, approved as low risk, sent for limited review, or escalated to full governance. List missing information, interim restrictions, and the decision deadline. Explain which reviewer is accountable, which are consulted, and who must be informed. Do not approve the use case; provide routing and rationale only. INTAKE: [PASTE].

LEADERSHIP CHECK *Your governance charter—not the AI output—determines authority.*

17 Pilot Charter

USE WHEN a use case is ready for controlled testing | **WORKSPACE** Word / project workspace | **RISK HIGH**

Have ready: *approved scope, owners, validation thresholds*

COPY / PASTE PROMPT

Draft a controlled pilot charter for [USE CASE]. Include objective, hypothesis, scope, exclusions, location, users, data, start and stop dates, accountable owner, clinical owner, technical owner, training prerequisites, baseline, success and balancing metrics, validation plan, human review, escalation, incident reporting, change control, communication, downtime, data retention, and explicit stop criteria. State what the pilot does not authorize. Add a weekly review agenda and a final decision template: stop, redesign, extend, or scale. Make all unknowns visible as bracketed decisions.

LEADERSHIP CHECK *A pilot is a bounded experiment, not an informal production launch.*

18 Monitoring and Drift Dashboard

USE WHEN a tool is entering ongoing operation | **WORKSPACE** Excel / BI / Word | **RISK HIGH**

Have ready: *baseline metrics, vendor monitoring, local logs*

COPY / PASTE PROMPT

Design an operating dashboard for [AI SOLUTION]. Include leading and lagging measures for input data quality, missingness, volume, latency, output distribution, confidence, override rate, false positives, false negatives, subgroup performance, exceptions, complaints, near-misses, incidents, downtime, adoption, manual workload, and realized benefit. Define owner, source, frequency, threshold, action, and escalation for each measure. Distinguish vendor-monitored from locally monitored signals. Add a version-change log and a rule that material changes trigger revalidation. Provide a monthly operating review agenda and a quarterly governance summary.

LEADERSHIP CHECK *If you cannot observe performance after go-live, you cannot claim the risk is controlled.*

19 Periodic Reapproval Review

USE WHEN conducting quarterly or annual governance review | **WORKSPACE** Word / GRC system | **RISK MODERATE**

Have ready: *original approval, performance data, changes, incidents*

COPY / PASTE PROMPT

Create a periodic reapproval packet for [AI USE CASE]. Compare current operation with the approved scope, data, model version, integrations, users, population, performance thresholds, and vendor terms. Summarize incidents, near-misses, complaints, overrides, drift signals, control failures, and realized value. Identify environmental changes such as formulary, policy, staffing, EHR configuration, law, or vendor ownership. Recommend renew, renew with conditions, revalidate, suspend, or retire, with rationale and owner. Include attestations for pharmacy, privacy, security, patient safety, and executive sponsorship.

LEADERSHIP CHECK *Reapproval should be evidence-based and time-bounded; "no complaints" is not performance evidence.*

20 Shadow AI Discovery Message

USE WHEN you need to surface unreviewed use without chilling reporting | **WORKSPACE** Outlook / Teams | **RISK** MODERATE

Have ready: *approved tools, reporting channel, amnesty approach*

COPY / PASTE PROMPT

Draft a message from pharmacy leadership inviting staff to disclose current or planned AI use. The tone should be curious, nonpunitive, and safety-focused. Explain that tools may already be embedded in common software; define examples of AI-assisted work; distinguish approved from reviewed; ask staff not to paste patient or confidential data into unapproved tools; provide one simple reporting channel; and offer office hours. State that near-misses and workarounds are valuable learning. Avoid fear, surveillance language, or a blanket ban. Provide a 120-word email, a 50-word Teams post, and manager talking points.

LEADERSHIP CHECK *Align the message with HR, compliance, and your actual policy before sending.*

MODULE 4

Build AI Literacy

Right-sized learning for staff, operators, governance leads, and executives

Literacy beats prohibition: people report risk sooner when they understand the boundary.

LEADERSHIP LENS Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

Three levels of literacy

AUDIENCE	CAPABILITY
All staff	What AI is/isn't; approved tools; data boundaries; verification; concern reporting.
Pharmacy operators	Tool SOP; output checks; override; downtime; documentation; escalation.
Governance leads	Risk assessment; evidence review; validation; monitoring; incident learning.

21 All-Staff AI Literacy Module

USE WHEN building a common baseline | **WORKSPACE** PowerPoint / LMS / Word | **RISK** MODERATE

Have ready: *approved-tool list, local policy, reporting channel*

COPY / PASTE PROMPT

Create a 20-minute AI literacy lesson for hospital pharmacy staff with no technical prerequisites. Cover: what AI is and is not; common pharmacy examples; why outputs can be fluent and wrong; approved tools and data boundaries; when verification is mandatory; how to document AI assistance; how to spot automation bias; how to report a concern or near-miss; and a two-minute scenario discussion. Use plain language, three memorable rules, five knowledge-check questions with answers, and a one-page job aid. Do not include vendor-specific promises or encourage clinical use outside approved workflows.

LEADERSHIP CHECK *Replace placeholders with local policy and require staff to know the reporting channel.*

22 Operator SOP Generator

USE WHEN staff need safe, repeatable steps for an approved tool | **WORKSPACE** Word | **RISK** HIGH

Have ready: *approved use, workflow, controls, screenshots if approved*

COPY / PASTE PROMPT

Draft a role-based SOP for [TOOL / USE CASE]. Include purpose, authorized users, prerequisites, allowed inputs, prohibited inputs, step-by-step workflow, required verification, documentation, override, escalation, downtime, incident reporting, version awareness, and competency check. Clearly separate system output from the pharmacist's accountable decision. Add examples of acceptable and unacceptable use and a "stop and escalate" box. Write for [ROLE] at an eighth-grade reading level without oversimplifying clinical or privacy controls. Mark every policy-dependent statement for local confirmation.

LEADERSHIP CHECK *Have the process owner test the SOP while performing the real workflow before release.*

23 Governance Lead Training Plan

USE WHEN committee members need deeper capability | **WORKSPACE** Word / LMS | **RISK** MODERATE

Have ready: *committee charter, review templates, local examples*

COPY / PASTE PROMPT

Design a six-session development path for pharmacy AI governance leads. Cover: AI lifecycle and roles; the five pharmacy risk categories; use-case scoping; vendor evidence; data governance; validation and human factors; monitoring and drift; incident response; procurement and exit; and communicating with P&T, compliance, legal, and the board. For each session, provide objectives, a pharmacy scenario, prework, exercise, decision artifact, facilitator profile, and competency measure. Finish with a capstone review of a fictional AI procurement and a rubric for consistent committee decisions.

LEADERSHIP CHECK *Train for decision quality and documentation—not technical vocabulary.*

24 Teach-Back and Competency Check

USE WHEN verifying that users can operate safely | **WORKSPACE** Forms / LMS | **RISK** HIGH

Have ready: *SOP, known failure modes, escalation path*

COPY / PASTE PROMPT

Create a competency assessment for users of [AI TOOL]. Include scenario-based questions that test intended use, prohibited use, PHI handling, output verification, uncertainty, override, downtime, near-miss reporting, and escalation. Add a teach-back script requiring the learner to explain in their own words what the tool cannot do and when they must stop. Provide scoring, remediation, and recheck criteria. Avoid trivia about how AI works; focus on safe decisions in the actual workflow.

LEADERSHIP CHECK *Completion is not competence. Include observed practice for high-risk workflows.*

25 Psychological Safety Conversation Guide

USE WHEN staff may hesitate to report AI concerns | **WORKSPACE** Team huddle / Word | **RISK** LOW

Have ready: *recent workflow changes, reporting mechanisms*

COPY / PASTE PROMPT

Create a 10-minute pharmacy team huddle that normalizes reporting AI errors, near-misses, confusing outputs, workarounds, and pressure to over-rely on automation. Include an opening statement from the leader, two realistic scenarios, five discussion questions, language for responding appreciatively to a report, and follow-up commitments. Emphasize that reporting protects patients and improves systems. Avoid blame, public identification, or promises that exceed the organization's just-culture policy.

LEADERSHIP CHECK *Leaders teach culture through their first reaction to bad news.*

26 Role-Based Training Matrix

USE WHEN coordinating different learning needs | **WORKSPACE** Excel / Word | **RISK** MODERATE

Have ready: *roles, use cases, risk tiers, annual training requirements*

COPY / PASTE PROMPT

Build a role-based AI training matrix for pharmacy staff, operators, informaticists, buyers, governance members, clinical leaders, and executives. Map each role to foundational knowledge, tool-specific skills, risk responsibilities, incident duties, competency method, refresher frequency, and evidence of completion. Distinguish awareness, working proficiency, and decision authority. Include triggers for retraining after material model changes, incidents, policy changes, or workflow redesign. Recommend the smallest practical curriculum that still matches each role's risk.

LEADERSHIP CHECK *Do not give low-frequency users more access than their training supports.*

MODULE 5

When Something Goes Wrong

Detect • Contain • Report • Investigate • Remediate and learn

Near-misses are free lessons. Capture them with the same discipline as errors.

> **LEADERSHIP LENS** Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

First-hour priorities

- Protect patients and maintain care continuity.
- Pause or constrain the tool; prevent additional exposure.
- Use the established reporting channel and name the incident lead.
- Preserve prompts, outputs, logs, model/version, timeline, and decisions.
- Communicate verified facts, unknowns, next decision, and next update time.

27 AI Incident Response Playbook

USE WHEN defining who does what before an event | **WORKSPACE** Word / incident system | **RISK** HIGH

Have ready: *existing safety, privacy, security, and downtime plans*

> **COPY / PASTE PROMPT**
> Draft an AI-related incident playbook that integrates with existing patient safety, privacy, security, compliance, and IT incident processes. Use five phases: detect, contain, report, investigate, remediate and learn. Define triggers, severity levels, on-call roles, decision authority, evidence to preserve, immediate containment options, patient-care continuity, notification questions for counsel, vendor coordination, communication, and return-to-service criteria. Cover wrong output, harmful recommendation, PHI exposure, bias, automation failure, unauthorized use, model change, and near-miss. Include a one-page first-hour checklist and a RACI.

LEADERSHIP CHECK *Do not create a parallel reporting universe; connect AI incidents to established channels.*

28 Tabletop Exercise Designer

USE WHEN practicing the response before a real event | **WORKSPACE** Facilitated workshop | **RISK** HIGH

Have ready: *playbook, participant roles, chosen use case*

COPY / PASTE PROMPT

Design a 60-minute tabletop exercise for this scenario: [DESCRIBE AI USE CASE]. The event begins with a plausible wrong output and unfolds through four injects: patient or operational impact; uncertain scope; vendor response; and external or executive attention. Provide facilitator notes, facts available at each stage, decisions expected, observer checklist, and debrief questions. Test detection, containment, reporting, evidence preservation, clinical continuity, communications, accountability, and return-to-service. Add success criteria and an after-action template. Do not reveal later injects to participants in advance.

LEADERSHIP CHECK *Exercise the real people named in the plan, including an executive decision-maker.*

29 Near-Miss Learning Review

USE WHEN an error was caught before harm | **WORKSPACE** Patient safety system / Word | **RISK HIGH**

Have ready: *de-identified event timeline, logs, workflow*

COPY / PASTE PROMPT

Conduct a just-culture near-miss review of the de-identified event below. Build a timeline and separate observable facts from assumptions. Analyze contributing factors across user interface, workflow, data, model behavior, training, staffing, environment, policy, handoffs, incentives, and vendor controls. Identify why the error occurred and why it was caught. Avoid blaming the last person in the chain. Recommend immediate containment, durable system changes, monitoring, owner, due date, and how to share the lesson without exposing protected or personnel information. EVENT: [PASTE DE-IDENTIFIED FACTS].

LEADERSHIP CHECK *Near-misses are high-value data. Track whether corrective actions actually close.*

30 First-Hour Executive Brief

USE WHEN leadership needs a fast, factual update | **WORKSPACE** Word / secure AI | **RISK HIGH**

Have ready: *verified facts, incident lead, current containment*

COPY / PASTE PROMPT

Draft a concise first-hour executive brief for an AI-related incident. Use only the verified facts provided. Include: what happened; when detected; known and unknown scope; patient, operational, privacy, security, and regulatory implications under assessment; current containment; care-continuity status; accountable incident lead; next decision; next update time; and requests for executive support. Label assumptions and unknowns. Avoid assigning blame, estimating liability, or making legal conclusions. FACTS: [PASTE APPROVED, DE-IDENTIFIED INCIDENT FACTS].

LEADERSHIP CHECK *Have the incident lead and counsel/compliance review before distribution.*

31 Root-Cause Question Set

USE WHEN the team is investigating beyond the symptom | **WORKSPACE** Word / investigation meeting | **RISK HIGH**

Have ready: *timeline, logs, change history, interviews*

COPY / PASTE PROMPT

Create an investigation plan for [INCIDENT]. Generate questions for users, workflow owners, IT, data teams, vendor, privacy, safety, and leadership. Examine intended use, scope drift, data quality, model version, configuration, integration, human factors, training, workload, alert design, overrides, access, logging, updates, and governance decisions. Request evidence that can confirm or disconfirm each hypothesis. Build a hypothesis tracker with supporting evidence, contradicting evidence, owner, and status. Distinguish root causes, contributing factors, and detection failures.

LEADERSHIP CHECK *Preserve logs and version information early; they may be ephemeral.*

32 Corrective Action and Learning Plan

USE WHEN moving from investigation to durable improvement | **WORKSPACE** Word / GRC system | **RISK HIGH**

Have ready: *root-cause findings, owners, due dates*

COPY / PASTE PROMPT

Convert the approved findings into a corrective and preventive action plan. For each cause or control failure, define immediate correction, durable system action, accountable owner, due date, evidence of completion, effectiveness measure, and review date. Prefer stronger controls—workflow redesign, technical constraint, automated detection—over reminders alone. Include SOP, training, vendor, contract, monitoring, and governance updates. Add criteria for safe return to service and a 30/60/90-day effectiveness check. Draft a concise lessons-learned summary for staff that supports transparency without exposing protected information.

LEADERSHIP CHECK *Close actions based on evidence of effectiveness, not merely task completion.*

MODULE 6

Your First 90 Days—and Beyond

Put governance to work, prove value, report clearly, and prepare for greater autonomy

Start with visibility and one accountable gate; mature through evidence, monitoring, and practice.

LEADERSHIP LENS Use AI to accelerate preparation, comparison, and documentation. Keep accountable people in control of clinical judgment, approvals, exceptions, and final decisions.

The 90-day sequence

WINDOW	OUTCOME
Days 1-30	Inventory AI; name owner; establish approval gate; publish interim rules.
Days 31-60	Run diligence; draft policy/charter; launch literacy; select controlled pilot.
Days 61-90	Drill incident response; activate monitoring; report to P&T/leadership.

33 Incident Trend Brief

USE WHEN reporting patterns to P&T or leadership | **WORKSPACE** Excel / Power BI / Word | **RISK** MODERATE

Have ready: *de-identified incident and near-miss data*

COPY / PASTE PROMPT

Analyze the de-identified AI incident and near-miss register for trends. Group by use case, risk category, severity, detection source, contributing factor, control failure, vendor, workflow stage, and time. Distinguish event count from exposure volume. Identify recurring patterns, weak signals, and areas where reporting may be low rather than risk absent. Produce a one-page leadership brief with three insights, three actions, owners, and measures. Do not infer causation from small counts or expose individual reporters.

LEADERSHIP CHECK *Normalize by use volume where possible; raw counts can mislead.*

34 Days 1-30 Launch Plan

USE WHEN starting governance with limited capacity | **WORKSPACE** Planner / Word | **RISK** MODERATE

Have ready: *current committees, approved tools, executive sponsor*

COPY / PASTE PROMPT

Create a practical 30-day launch plan for pharmacy AI governance. Priorities: name an accountable owner and sponsor; inventory AI already in use; establish one intake and approval gate; publish interim data and use rules; identify high-risk active uses; and define the incident reporting path. Use weekly milestones, named roles, deliverables, decision points, and a simple measure of completion. Reuse existing committees and processes where possible. Include a communication plan that invites disclosure rather than punishment. Keep the plan achievable without a dedicated AI department.

LEADERSHIP CHECK *The first month should create visibility and control, not a perfect policy.*

35 Days 31-60 Build Plan

USE WHEN moving from discovery to repeatable governance | **WORKSPACE** Planner / Word | **RISK** MODERATE

Have ready: *inventory, draft charter, vendor pipeline*

COPY / PASTE PROMPT

Create a days 31-60 implementation plan. Priorities: complete risk triage of active uses; standardize vendor due diligence; draft the governance charter and review matrix; define pilot and monitoring templates; launch foundational AI literacy; and select one low-to-moderate-risk pilot. Include owners, dependencies, evidence of completion, and unresolved executive decisions. Add a weekly operating rhythm and a short P&T/leadership update.

LEADERSHIP CHECK *Choose one pilot that builds governance capability, not only excitement.*

36 Days 61-90 Operationalize Plan

USE WHEN turning the framework into an operating system | **WORKSPACE** Planner / Word | **RISK** MODERATE

Have ready: *approved pilots, incident playbook, metrics*

COPY / PASTE PROMPT

Create a days 61-90 plan. Priorities: run an AI incident tabletop; activate monitoring for approved tools; complete role-based training; review vendor and contract gaps; establish periodic reapproval; report status and risks to P&T and executive leadership; and publish the next-quarter roadmap. Include a maturity score at day 90, evidence for each rating, and the three highest-value next investments.

LEADERSHIP CHECK *Day 90 is the beginning of a cadence, not the end of governance.*

37 P&T / Executive Dashboard

USE WHEN reporting progress without technical overload | **WORKSPACE** PowerPoint / Word | **RISK** MODERATE

Have ready: *inventory, risk register, pilot results, incidents, value metrics*

COPY / PASTE PROMPT

Design a one-page AI governance dashboard for P&T and executive leadership. Include: active use cases by risk tier and status; decisions needed; patient-safety and privacy controls; validation status; incidents and near-misses; training coverage; vendor issues; realized value; and next 90-day priorities. Use plain language and show trend, threshold, owner, and action—not decorative metrics. Add a short narrative explaining where leadership attention is required and what would trigger suspension of a tool.

LEADERSHIP CHECK *A dashboard should support decisions; omit metrics that do not change action.*

38 AI Portfolio Prioritization

USE WHEN choosing where to invest next | **WORKSPACE** Excel / strategy workshop | **RISK** MODERATE

Have ready: *candidate use cases, strategic goals, constraints*

COPY / PASTE PROMPT

Rank candidate pharmacy AI use cases using value, feasibility, readiness, evidence, risk, workflow burden, data availability, integration, change capacity, and reversibility. Separate patient-care, operational, administrative, and strategic uses. Apply a penalty for unclear ownership, weak baselines, inaccessible data, or irreversible automation. Provide a portfolio map with quick wins, capability builders, strategic bets, and avoid/defer. Recommend a balanced 12-month sequence and explain what must become true before deferred use cases can proceed.

LEADERSHIP CHECK *Prioritize problems worth solving, not technologies looking for a home.*

39 Business Case and Total Cost

USE WHEN testing whether projected ROI is real | **WORKSPACE** Excel / Word | **RISK** MODERATE

Have ready: *baseline volume, labor, quality, implementation, licensing*

COPY / PASTE PROMPT

Build a business-case model for [SOLUTION]. Include baseline process cost and performance; one-time implementation, integration, validation, training, change, legal, security, and data costs; recurring licensing, support, monitoring, model updates, audit, and governance costs; productivity ramp; downtime and fallback; expected benefit; confidence ranges; and opportunity cost. Separate cash savings, cost avoidance, capacity, quality, experience, and risk reduction. Create conservative, expected, and optimistic scenarios, identify assumptions with the greatest sensitivity, and define measures needed to prove value after launch.

LEADERSHIP CHECK *Do not convert theoretical time saved into cash savings unless capacity or expense actually changes.*

40 Implementation Timeline Challenge

USE WHEN a vendor timeline feels aggressive | **WORKSPACE** Word / project plan | **RISK** MODERATE

Have ready: *vendor plan, interface needs, local approvals*

COPY / PASTE PROMPT

Stress-test the proposed implementation timeline for [SOLUTION]. Identify dependencies for contracting, BAA, security, interfaces, data mapping, validation, workflow design, training, change management, support, downtime, measurement, and governance approval. Compare vendor tasks with customer tasks and name the critical path. Highlight assumptions, resource conflicts, likely rework, and decisions needed by date. Produce a realistic range rather than a single date and define conditions for a safe limited pilot versus production scale.

LEADERSHIP CHECK *An aggressive timeline often hides customer work; make every dependency visible.*

41 Workflow Opportunity Scan

USE WHEN looking for high-value use cases before shopping | **WORKSPACE** Workshop / Word | **RISK** MODERATE

Have ready: *workflow map, pain points, volume and quality data*

COPY / PASTE PROMPT

Analyze the pharmacy workflow below to identify where AI, conventional automation, process redesign, or no technology is the best intervention. Look for high-volume cognitive work, repeated summarization, classification, forecasting, queue prioritization, documentation, inventory, scheduling, and decision support. For each opportunity, describe the problem, baseline, affected roles, data, benefit, risk, simpler alternatives, and first experiment. Explicitly identify steps that should be eliminated or standardized before automation. Rank opportunities by value, feasibility, and safety.

LEADERSHIP CHECK *Redesign the workflow before digitizing a workaround.*

42 Agentic AI Readiness Review

USE WHEN a vendor proposes software that takes multi-step action | **WORKSPACE** Word / governance review | **RISK** HIGH

Have ready: *agent capabilities, tools/actions, permissions, environment*

COPY / PASTE PROMPT

Assess the readiness of [AGENTIC AI USE CASE]. Map every action the agent can take, system it can access, permission it holds, data it can read or write, decision it can influence, and downstream effect. Identify boundaries, human confirmation points, least-privilege controls, transaction limits, sandbox needs, observability, rollback, kill switch, credential management, memory and retention, prompt-injection exposure, tool failure, and cascading-error scenarios. Recommend the maximum safe autonomy level for a pilot and the evidence required to increase it. Assume that multi-step errors can compound.

LEADERSHIP CHECK *Autonomy should expand only after evidence; start with read-only or draft-only where feasible.*

43 Board-Level AI Risk Brief

USE WHEN the board or C-suite asks how the organization is protected | **WORKSPACE** PowerPoint / Word | **RISK** MODERATE

Have ready: *governance status, material risks, metrics, decisions*

COPY / PASTE PROMPT

Draft a board-ready briefing titled “How Pharmacy Is Governing AI Risk.” Explain in plain language: where AI is used; the five pharmacy risk categories; one approval gate; vendor and validation controls; human accountability; monitoring; literacy; incident readiness; and the 90-day roadmap. Include current strengths, material gaps, risk acceptance decisions, investment needs, and three questions for board oversight. Avoid technical detail, fear, and unsupported assurances. Use evidence and distinguish control design from control effectiveness.

LEADERSHIP CHECK *Never say “zero risk.” Explain what is known, controlled, monitored, and accepted.*

44 Peer Discussion Facilitator

USE WHEN turning conference learning into local action | **WORKSPACE** Team workshop | **RISK** LOW

Have ready: *participant roles, local priorities*

COPY / PASTE PROMPT

Facilitate a 30-minute pharmacy leadership discussion using these questions: Which AI tools are already in use that never went through review? Who owns the outcome when an AI-assisted decision goes wrong? What would our first incident-response drill look like? Add six follow-up questions on vendor evidence, data boundaries, workflow fit, literacy, near-miss reporting, and the first 90 days. Provide a safe opening, time boxes, capture template, dot-voting method, and a closing commitment with owner and date.

LEADERSHIP CHECK *End with one accountable next step, not a list of interesting observations.*

REFERENCE

Standards compass for pharmacy AI buyers

Use these sources as a starting compass, then confirm current requirements with your organization's legal, compliance, privacy, security, safety, informatics, and pharmacy leaders.

SOURCE	WHY IT MATTERS	LINK
NIST AI RMF 1.0 and Generative AI Profile	Voluntary risk-management structure: Govern, Map, Measure, Manage; includes generative-AI risk actions.	https://www.nist.gov/itl/ai-risk-management-framework
HHS OCR HIPAA Security Risk Analysis Guidance	Risk analysis should cover confidentiality, integrity, and availability of all ePHI the organization creates, receives, maintains, or transmits.	https://www.hhs.gov/hipaa/for-professionals/security/guidance/guidance-risk-analysis/index.html
ONC HTI-1 Algorithm Transparency	Transparency expectations for predictive decision support interventions in certified health IT, including intended use, populations, limitations, and risk management.	https://healthit.gov/regulations/hti-rules/hti-1-final-rule/
FDA Digital Health Guidance Navigator	Current status of device-software, AI-enabled device, clinical decision support, cybersecurity, and change-control guidance.	https://www.fda.gov/medical-devices/regulatory-accelerator/medical-device-software-guidance-navigator
ASHP Artificial Intelligence Resource Center	Pharmacy-specific statements, guidance, and practice resources for understanding and evaluating AI applications.	https://www.ashp.org/pharmacy-practice/resource-centers/informatics/artificial-intelligence

What "good" looks like at the end of 90 days

- A complete AI inventory with owners, risk tiers, approval status, and next review date.
- One intake and approval gate connected to existing pharmacy and enterprise governance.
- Evidence-based vendor diligence and contract issue lists for active procurements.
- At least one controlled pilot with baseline, thresholds, monitoring, and stop criteria.
- Role-based training with competency checks and a known concern-reporting channel.
- An AI incident playbook tested through a tabletop exercise.
- A concise P&T/executive dashboard showing risk, value, decisions, and next actions.

Leader’s field notes

Use this page after the session or during your first governance meeting.

The AI tool already in use that most needs review:

The person who should own pharmacy AI governance:

The first vendor question we need to start asking:

The workflow where AI could create the most value:

The workflow where AI could create the most harm:

The first incident scenario we should drill:

My next action, owner, and date:

THREE DISCUSSION STARTERS Which AI tools are already in use that never went through review? Who owns the outcome when an AI-assisted decision goes wrong? What would your first incident-response drill look like?

About the author

Paul J. Swider is CEO and Chief AI Officer of RealActivity. He brings decades of healthcare IT and operations experience to practical AI implementation, governance, and risk management. His approach emphasizes partnership over sales, proven outcomes over promises, and guardrails that help leaders adopt AI with confidence.

> **SESSION COMPANION** AI Risk Management for Pharmacy Leaders: Building Guardrails Before You Need Them — Hospital Pharmacy RxBuyer Conference, Orlando, August 10, 2026.

Educational and legal notice

This guide is for general educational and operational planning purposes. It is not legal, regulatory, medical, cybersecurity, contracting, or professional advice. AI outputs can be inaccurate, incomplete, biased, outdated, or inappropriate for a particular patient, workflow, organization, or jurisdiction. Do not enter PHI, credentials, confidential information, or proprietary data into tools that your organization has not approved for that purpose. Maintain qualified human review for clinical and other high-impact decisions. Confirm current law, regulation, policy, professional standards, product status, and contract requirements with the appropriate qualified reviewers.

Prepared July 29, 2026. Standards and product capabilities change; verify current source material before use.