International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 1125-1133
http://www.ijcesen.com

ISSN: 2149-9144

Research Article

A Zero-Trust Compliance Architecture for LLM-Integrated Pharmaceutical IT Systems: Securing AI-Assisted Workflows with Data Integrity and Regulatory Controls

Sreeharsha Amarnath Rongala*

Epic Pharma LLC, USA
* **Corresponding Author Email:** rongala.rsa@gmail.com - **ORCID:** 0000-0002-5247-0150

Article Info:

DOI: 10.22399/ijcesen.497
Received : 29 December 2025
Revised : 20 February 2026
Accepted : 22 February 2026

Keywords

Zero Trust LLM Security,
Pharmaceutical AI Compliance,
GxP AI Integration,
LLM Data Integrity ALCOA,
AI Governance Pharmaceutical Quality

Abstract:

LLMs are being used in regulated pharmaceutical IT workflows for standard operating procedure (SOP) generation, deviation triage and Corrective and Preventive Action (CAPA) writing, but this is complicated by stochastic responses and susceptibility to prompt injection along with issues of data integrity, role-based access control and auditability. This paper proposes the Zero-Trust Compliance Architecture (ZT-LLM-COMPASS) for this zero-trust setting, where every model invocation, context retrieval, tool invocation, and artifact generation is treated as a potential attacker payload. ZT-LLM-COMPASS incorporates four LLM Security Planes: fine-grained authorization via Identity & Policy Enforcement Plane, prompt firewalls and controlled retrieval via Prompt & Context Security Plane, least-privileged permissions on function calls via Tool Sandbox, and tamper-obvious audit packages via Evidence Plane. Human-in-the-loop verification gates ensure that the LLM can only propose and never commit regulated records. Adversarial evaluation demonstrates that the design is strong against injection attacks, leakage, and auditing challenges, enabling high AI productivity while retaining regulatory compliance.

1. Introduction

1.1 Problem Statement: LLMs in Regulated Documentation and Decisions

LLMs are being used to draft SOPs, triage deviations, and create CAPAs in quality processes in the pharmaceutical industry. While they may increase efficiency, permanent documentation in a regulated environment poses challenges. In addition, the stochastic nature of transformer models means wide-ranging validation to ensure that the model behaves as expected. Indirect prompt injection attacks can manipulate the transformer behavior without prompt injections directly into the target model. Outputs of LLMs will be controlled records subject to regulation, but their provenance (training data, retrieved documents, or injection) will remain opaque.

1.2 Why Classic "LLM Security" is Insufficient for GxP

Current LLM security emphasizes perimeter defenses and input sanitization, but GxP environments demand ALCOA+ principles: accuracy, legibility, contemporaneousness, originality, and attributability. LLM-generated content must carry verifiable provenance showing which controlled templates informed it, who reviewed it, and when modifications occurred. Zero Trust frameworks designed for microservices don't address generative model risks: context poisoning through retrieval manipulation, tool abuse via function-calling, and auditing probabilistic outputs. Pharmaceutical organizations must architect verifiable evidence chains satisfying both cybersecurity controls and regulatory expectations for electronic records.

1.3 Paper Contributions

This article presents the Zero-Trust Compliance Architecture (ZT-LLM-COMPASS) with four enforcement planes, a control-objective expert-based mapping of GxP regulatory clauses to auditable compliance methods, a validation

methodology leveraging adversarial testing for injection resistance and audit completeness, and human-in-the-loop verification gates as LLM implementation patterns for proposing (but never committing) regulated records.

1.4 Scope Boundaries

This research addresses LLM-assisted workflows where LLMs augment human decision-making, with tasks requiring LLMs including authoring documents, retrieving data from controlled databases and repositories, and invoking limited tools through a supervised broker. The scope excludes workflows involving the training of proprietary foundation model architectures, as well as workflows where agents independently make final decisions in regulated processes without human supervision.

2. Background and Related Work

2.1 Zero Trust Principles Applied to AI/LLM Systems

Zero Trust architecture requires identity, device posture, and contextual risk to be evaluated per transaction. LLM systems extend the definition of transaction and principals. When prompts are sent, they must be authenticated for the user, the context of the workflow, and the sources retrieved. The LLM as a service principal is assigned scoped permissions according to least privilege principles which limits its access. Based on scenarios that comprise applicable policy decision points, users may be blocked at various points in a workflow, and in various LLM sessions, from accessing certain sections of various controlled documents for certain applications. Continuous verification assumes model outputs are untrustworthy until independently verified against business rules and regulatory constraints.

2.2 GxP Computerized Systems Fundamentals: Audit Trails, Electronic Records, Data Integrity

GxP requires computerized systems to provide a full audit trail describing who did what, when and why. Electronic records must comply with ALCOA+ (attributed, legible, contemporaneous, original, and accurate) and additional requirements. Human-reviewed and approved versions of electronic records must be retained by the systems, including prompts, versioned retrieved contexts, model versions and drafts, and outputs for human review. Controls include preventive controls (e.g., no access to the model), detective controls (e.g., detection of model output anomalies), and corrective controls (e.g., reconstructing audit trail). The missing bridge between current LLM security and GxP requirements is validated evidence chains that inspectors can review.

2.3 LLM Risk Landscape: Prompt Injection, Context Poisoning, Data Leakage, Model Drift, Tool Abuse

Prompt injection attacks rely on the inability to distinguish between instructions and data in natural language interfaces. An indirect prompt injection attack puts malicious instructions within the retrieved documents, while context poisoning attacks bias either the retrieval rank or the embedding that the LLM considers authoritative. Data leakage may occur from either memorization in LLMs or from specific prompting. Model drift may occur through provider updates or fine-tuning. Tool abuse occurs when function calling LLMs are allowed to make expensive database writes or arbitrary API calls. Pharmaceutical manufacturing compounds the tool abuse risk, corrupting quality records and exposing the company to regulatory scrutiny.

2.4 Gap: A Compliance-Evidence Architecture Tying LLM Controls to Inspection-Ready Artifacts

Existing LLM security frameworks can address many technical vulnerabilities but do not track the underlying provenance explaining how particular SOP phrases were generated. Compliance frameworks offer controls that are inadequate or too vague for probabilistic systems. The biggest gap it addresses is providing evidence chains that meet both cybersecurity and GxP requirements. ZT-LLM-COMPASS articulates control objectives in regulatory terms, maps them to controls and technical implementation and describes evidence packages that may be used to show compliance at an inspection.

3. Regulated Use Cases and Compliance Requirements

3.1 Use Case A: SOP Drafting Assistant

SOP drafting assistants help write new procedures, for example, by retrieving a template and letting the user select compliant language. The user chooses a procedure type and applicable regulatory scope. The LLM retrieves controlled template sections, finds associated procedures, and produces draft text containing the required elements (purpose

statements, responsibility matrices, etc.). The output is generated text, which is marked as draft content versus controlled documents.

3.2 Use Case B: Deviation Triage Assistant

Deviation triage assistants help subject matter experts review quality events, with functions such as reading a deviation's description, acquiring previous deviations, and recommending classifications. They have read-only access to the deviation histories but cannot modify controlled data. It suggests initial severity classifications and rationale, and generates customized lists of questions for further investigation. These require independent verification by an expert.

3.3 Use Case C: CAPA Drafting Assistant

CAPA drafting assistants allow quality users to draft actions and timelines based on investigation findings, apply similar historical CAPAs, and suggest checks to determine effectiveness. Assistants cannot assign actions or timelines. They cannot close CAPA records nor sign off on effectiveness checks. Approval from authorized personnel is required to execute any tool commands that modify controlled records.

3.4 "What Becomes a Record?" Decision Points

Pharmaceutical information management systems distinguish between working documents and controlled records. Draft human readable and machine readable outputs from LLMs do not have regulated status until manual human review and approval workflow and quality management system records are in place. Drafts are held in intermediate environments with relaxed change control, while approved content is held in immutable controlled repositories with formal change control.

3.5 Compliance Requirements Distilled into System Requirements

These principles map onto technical controls for identity, provenance, knowledge sources, auditability, retention and segregation of duties. Identity controls afford qualified personnel access to workflow features. Provenance mechanisms track all steps in the workflow, from prompt, through model inference, to approval. Knowledge sources control access so that only approved sources are queried. Auditability for prompts, retrievals and approvals is realized via tamper-obvious logs. Retention policies determine how long these evidence packages exist. Segregation of duties prevents draft generators from approving their own content as controlled records.

4. Threat Model Re-Framed as Data-Integrity Failure Modes

4.1 Assets: Controlled Documents, Quality Records, Audit Trail, Knowledge Base, Identities/Roles, Tool Endpoints

Critical assets include controlled documents, quality records, audit trails, knowledge bases, user identities, and tool endpoints. Controlled documents are definitive sources of knowledge. Quality records are permanent records for compliance. Audit trails record decision-making. Knowledge bases include metadata, embeddings, and retrieval indices. User identities authenticate users. Tool endpoints allow privileged function calling.

4.2 Adversaries: Malicious Insider, Compromised Account, External Attacker, Accidental Misuse

Attackers include rogue insiders working with legitimate credentials, account compromise giving external attackers access from the inside, attacks via public interfaces, supply chain attacks through shared document repositories, and user interfaces and training sessions that confuse users into unintentionally making mistakes that may have damaging consequences when using the technology.

4.3 Abuse Cases Mapped to SOP/Deviation/CAPA Workflows

Prompt injection in SOP writing might be embedded into the instruction to omit specific steps. Retrieval manipulation in deviation triages could result in only finding superseded procedures and incorrect severity levels. CAPA tool abuse can include unauthorized premature closure of corrective actions. Data exfiltration exploits use the text generation method to encode sensitive data for exfiltration. Model updates introduce validated-state drift, invalidating prior validation evidence.

4.4 Risk Classification and Why Continuous Verification is Required

The highest impact issue could be record corruption leading to non-conformances that can only be detected during inspection, exposing compliance risk and potential recalls. Likelihood may differ from customary software because adversarial input

can exploit emergent behaviors of models. The attack surface continuously changes with each model update, knowledge base update, and workflow change and should be kept under control throughout the system lifecycle via continuous monitoring and periodic adversarial testing. Table 1 categorizes the primary risk vectors associated with LLM integration in pharmaceutical IT systems, mapping each threat type to its specific impact on GxP compliance requirements and the affected regulatory principle (ALCOA+).

5. ZT-LLM-COMPASS: Zero-Trust Compliance Architecture

5.1 Design Goals: Inspection Readiness, Least Privilege, Verifiable Provenance, Controlled Record Lifecycle

ZT-LLM-COMPASS achieves four properties: producing auditor-interpretable evidence of inspection as to whether the component is ready to be inspected, least-privileged components, verifiable provenance from user intent to approved output, and well-defined version history supporting draft content, pending output and approved records separately.

5.2 Trust Boundaries and System Components

The architecture has interfaces in three zones, protecting trust boundaries between user devices in untrusted zones that provide prompts to authenticated channels, LLM inference services in semi-trusted zones that make decisions under policy constraints, and controlled document repositories in trusted zones that are accessed through policy-mediated interfaces that validate every transaction.

5.3 Plane A: Identity & Policy Enforcement

The Identity & Policy Enforcement Plane specifies what actions can occur under given circumstances. Identity systems using identity-based access control bind verifiable credentials to users to reflect device posture and network location. Attribute-based access control policies can express authorizations at a fine granularity, such as users' roles, document categories, workflow status, and times of the day. They can be version-controlled and automatically tested with policy-as-code. Decision logging captures the result of every policy evaluation.

5.4 Plane B: Prompt & Context Security

The Prompt and Context Security Plane controls what information flows into and out of LLMs. Prompt firewalls filter prompts looking for, e.g., injection patterns or policies. In data-minimization, retrievals are limited to the minimum document fragment for the task, and controlled RAG uses-only approved, version-controlled repositories. Provenance tracking identifies retrieved documents and sections, storing version IDs and associating them with generated text.

5.5 Plane C: Tool Sandbox

Tool Sandbox enables LLMs to run against enterprise systems and APIs. Brokers allow for the enforcement of allowlists, parameter limits, and different privileges. By default, draft workflows are read-only, separating reading from writing. Write operations enforce step-up approval with user review and electronic signature. Allowlists restrict available functionality per workflow type. Parameter constraints enforce function argument validation to ensure compliance with both type and business rule requirements.

5.6 Plane D: Evidence & Integrity

The Evidence & Integrity Plane composes tamper-clear audit packages for each interaction that include the prompts, context, model choices, and human decisions. These structured artifacts include the prompt text, retrieved documents, model identifiers, non-volatile full outputs (including rejected completions), and human decision timestamps. Append-only logging preserves contemporaneous records. Cryptographic signatures detect tampering. Exports generate inspector-readable reports, and retention policies store audit packages for controlled record lifecycles. Table 2 details the four security planes of the ZT-LLM-COMPASS architecture, describing the primary function of each plane, key technical mechanisms implemented, and the trust boundary enforced.

6. Workflow Implementations

6.1 SOP Drafting Workflow

Authorized users initiate drafting of SOPs by submitting requests, identified by document category and regulatory scope. The system validates qualifications and provisions for temporary drafting workspaces. For document retrieval, approved templates are selected with a constraint on the sections to be returned. The LLM receives prompts and controlled sections of the document. Tool Sandbox contains read-only

functions and operates as a human-in-the-loop gatekeeper, including checks for accuracy, completeness, and compliance with regulations and policies. After approval, Evidence Plane archives the audit packages and exports them to the change control workflows.

6.2 Deviation Triage Workflow

This process is based on archiving read-only copies of descriptions of events and their containment procedures, and can only retrieve historical deviations of the event type. The LLM generates initial classifications with justifications. It suggests questions for investigation, stating these are merely suggestions to be verified by experts. While read-only queries are possible with Tool Sandbox, write queries are prohibited. The Evidence Plane logs all activities, including justification of classifications and experts' final decisions.

6.3 CAPA Drafting Workflow

CAPA drafting retrieves previous CAPAs with similar roots and reasons, and generates action items to eliminate the root cause. Hypothesis labeling is used to label all outputs as such. Action constraints prevent plans from violating business rules. Tool constraints are highly limited compared to other workflows. Multiple approval steps form an approval chain, and Evidence Plane records reviewer disposition comments at each step. Table 3 compares the three primary pharmaceutical quality workflows (SOP drafting, deviation triage, and CAPA drafting) in terms of LLM capabilities, human-in-the-loop requirements, and tool access restrictions under the ZT-LLM-COMPASS architecture.

7. Control Objectives Mapped to Mechanisms and Evidence

7.1 Control-Objective Framework

The control-objective framework states regulatory objectives in terms of testable system attributes. Each objective specifies what the system is expected to do, relevant criteria for success, and what evidence is needed. Control objectives are generally hierarchical. Higher-level objectives decompose into specific control mechanisms. For preventive controls, failure to let through test probes is acceptable negative testing. For detective controls, timely identification of test probes is a successful test. Corrective controls are tested by incident response exercises.

7.2 Mapping Table: Control Objective → Mechanism → Evidence for SOP/Deviation/CAPA

The mapping table offers traceability from compliance requirements through mechanisms to verification evidence. For example, the requirement "prevent unauthorized SOP modifications" maps to the mechanisms "role-based access control with workflow-context evaluation", and verification evidence includes "access logs showing denied attempts, role assignment records, and audit trail entries". For deviation triage, "ensure classification decisions are attributable to qualified personnel" would mean "require SME explicit approval before commits" and "classification logs with LLM suggestions, SME approval signatures, and divergence tracking".

7.3 How This Mapping Supports Audit/Inspection Narratives

By mapping control objectives to navigable inspection responses, quality teams can show satisfactory compliance with the Prompt & Context Security Plane controls, showing detection logic, validation results and sample audit packages for attempts that are blocked when they ask the system how it is protected against prompt injection, as well as which evidence artifacts correspond with which regulatory requirements.

8. Validation, Monitoring, and Change Control

8.1 Intended Use and Risk Assessment Per Workflow

Validation starts with the intended use, followed by drafting SOPs (Standard Operating Procedures) for approval before controlled use. In deviation triage, decisions wait for expert review rather than being validated or invalidated. The CAPA drafting process includes action planning requiring approval from quality management. Risk assessment techniques measure processes against threat model failure modes. Risk mitigation includes controlled RAG governance, read-only access to data, and approval chains.

8.2 Validation Strategy: Requirements, Tests, Evidence Packages

Validation informs requirements for an intended use, test procedures to validate requirements for benign and malicious scenarios, and evidence packages of test operations and results. Requirements include measurable capabilities such

as a detection rate for injection patterns. These tests include functional testing and adversarial testing, and the evidence packages consist of test plans, test execution logs, comparison analysis, and the root cause analysis. The adversarial tests provide values for resistance to injection, false positive rate, leakage prevention and audit completeness.

8.3 Operational Monitoring: Anomaly Detection, Exception Handling, Periodic Review

Runtime issue alerting (indications of threats or degradation) and anomaly detection (recognizing patterns in prompts, retrieval frequency, and outputs). Exception handling establishes how to respond when monitoring indicates an issue. Periodic reviews of monitoring results allow detection of trends, evaluation of control effectiveness, and improvement.

8.4 Re-Qualification Triggers and Lifecycle Controls

All changes affecting output behavior or security controls (e.g. of a model, version or a system) that could change the validated behavior trigger re-qualification (i.e. verification with validation testing). Knowledge base updates trigger verification of retrieval mechanism versioning. The degree of such re-testing is risk based, corresponding to the change impact level. Change control procedures ensure that changes to production environments are reviewed, tested, and quality assured before deployment. Table 4 outlines the validation and continuous monitoring controls for ZT-LLM-COMPASS, specifying the control type, validation approach, monitoring mechanism, and re-qualification trigger for each category.

9. Discussion

9.1 Tradeoffs: Usability vs Control Friction, Speed vs Evidence Rigor, Explainability Limits

Thorough controls introduce friction to user experience and workflow speed, while prompt firewalls induce latency in content analysis. Constrained retrieval may miss useful sources. Tool-sandboxing requires additional steps, and explicit approval. Generally, there is a trade-off between the speed and rigor of generated evidence, and the whole package: generating an audit package requires considerable process and storage resources. Explainability is difficult because while it is easier to explain which docs were retrieved, it is harder to explain generated phrasing.

9.2 Implementation Patterns: Phased Rollout, Draft-Only Mode First, Gradual Tool Enablement

Applications will gradually roll out more LLM capabilities with more security. For example, early implementations may be draft-only, where outputs are advisory (i.e., may be wrong but provide useful inputs for the user), with permissions expanded as internal evidence grows. Using a staged rollout and active validation, as controls are added, user feedback is used to improve the interface and ensure that the controls improve rather than obstruct quality processes.

9.3 Generalization Beyond SOP/Deviation/CAPA

Other pharmaceutical IT processes that could apply ZT-LLM-COMPASS principles include change control workflows (for example, to assist with impact assessments) and investigation management workflows that employ triage methods to look across different sites for patterns of non-conformance. Training management systems using LLM can be used, but each use case needs to specialize their four planes according to the workflow risks. Certain patterns apply, including treating outputs as untrusted until verified, limiting retrieval to trusted sources, sandboxing tools' access, and generating audit packages.

10. Limitations and Future Work

Current implementations are limited, and the variation in quality system designs across multiple sites complicates architectural design. The opacity of LLMs prevents verification of their reasoning processes, forcing the field to rely on empirical evaluation. Future work might include audit package schemas for standardization, adversarial test corpora for LLMs in pharmaceutical environments, cross-organizational metrics to assess the evidence of control effectiveness, explainable AI methods for validation in compliance scenarios, and standardization of common adversarial techniques applicable to different datasets for LLMs. Creating these elements would support the construction of a framework, guiding LLM development processes and enabling consistent compliance across organizations.

Table 1: *LLM Risk Categories and GxP Compliance Impacts*

Risk Category	GxP Compliance Impact	Affected ALCOA+ Principle
Prompt Injection Attacks	Manipulation of LLM behavior leading to generation of non-compliant procedure text or omission of required regulatory steps	Accuracy, Originality
Context Poisoning via Retrieval Manipulation	Biased outputs from outdated or superseded controlled documents resulting in incorrect quality decisions	Accuracy, Contemporaneousness
Data Leakage through Model Memorization	Exposure of proprietary formulations or confidential compliance information violating data protection requirements	Originality, Attributability
Model Drift from Provider Updates	Invalidation of validated-state evidence requiring system re-qualification and potential regulatory reporting	Consistency, Completeness
Tool Abuse via Function Calling	Unauthorized modification of controlled records bypassing approval workflows and audit trail requirements	Attributability, Enduring

Table 2: *ZT-LLM-COMPASS Four-Plane Architecture Components*

Architecture Plane	Primary Function	Key Technical Mechanisms
Identity & Policy Enforcement Plane	Authenticate users and authorize actions based on workflow context, roles, and device posture	Strong identity binding, Attribute-Based Access Control, Policy-as-code implementation, Decision logging
Prompt & Context Security Plane	Control information flow into LLMs through filtering, validation, and constrained retrieval from approved sources	Prompt firewall with injection detection, Data minimization, Controlled RAG governance, Provenance tracking
Tool Sandbox Plane	Mediate LLM interactions with enterprise systems through allowlists, parameter validation, and privilege separation	Function call broker, Read/write separation, Step-up approval requirements, Parameter constraint enforcement
Evidence & Integrity Plane	Generate tamper-evident audit packages linking prompts, context, model behavior, and human decisions	Append-only logging, Cryptographic signatures, Inspector-readable exports, Retention policy enforcement

Table 3: *Pharmaceutical Workflow Implementation Characteristics*

Workflow Type	LLM Capabilities Permitted	Human-in-the-Loop Verification Requirements
SOP Drafting Assistant	Generate draft procedure text from controlled templates, Retrieve related procedures, Suggest compliant language patterns	Explicit approval before submission to change control, Verification of accuracy and regulatory alignment, Electronic signature required
Deviation Triage Assistant	Analyze event descriptions against historical patterns, Suggest preliminary severity classifications, Generate investigation question lists	Subject Matter Expert verification required before classification commits, Final decision authority remains with qualified personnel, Divergence tracking when overriding LLM suggestions
CAPA Drafting Assistant	Propose corrective action plans based on root cause analysis, Retrieve similar historical CAPAs, Suggest effectiveness verification methods	Quality management approval for all action assignments, Multiple approval stages with segregation of duties, Cannot close records or verify effectiveness checks
Common Constraints Across All Workflows	Read-only access to controlled repositories during drafting phase, Hypothesis labeling for all generated outputs, Citation of source documents	Draft content remains in temporary workspaces until human approval, Approval workflows enforce segregation of duties, Evidence packages sealed before controlled record commitment

Table 4: *Validation and Monitoring Control Framework*

Control Category	Validation Approach	Continuous Monitoring Mechanism

Injection Resistance	Adversarial testing with curated prompt corpus containing known injection patterns and embedded instructions	Anomaly detection on prompt patterns, Firewall detection rate tracking, Periodic adversarial test re-runs
Retrieval Authorization	Functional testing attempting access to unauthorized repositories, draft documents, and obsolete versions	Authorization decision logging, Retrieval source verification, Version control compliance checks
Tool Sandbox Enforcement	Negative testing of unauthorized write operations, parameter boundary violations, and privilege escalation attempts	Tool invocation monitoring, Parameter validation failure tracking, Step-up approval enforcement verification
Audit Trail Completeness	Evidence package schema validation, Completeness checks for all required elements, Hash chain verification	Audit package generation monitoring, Missing element detection, Cryptographic signature validation
Human-in-the-Loop Gate Integrity	Testing direct commit attempts bypassing approval workflows, Segregation of duties violation attempts	Approval workflow enforcement tracking, Electronic signature capture verification, Blocked commit attempt logging

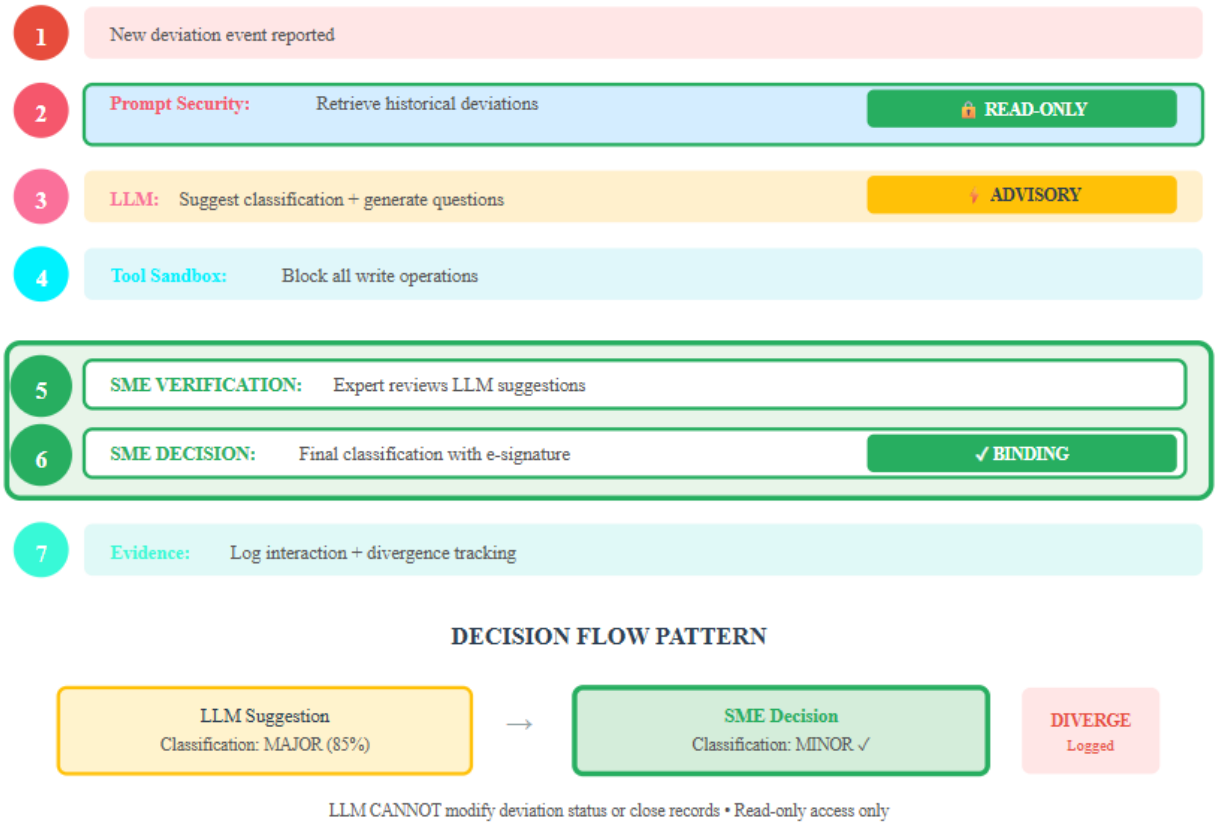

Figure 1: *Deviation Triage with Read-Only Context and SME Verification*

11. Conclusions

Adapting large language models (LLMs) to pharmaceutical quality systems could bring efficiency gains. However, LLMs introduce data integrity, audit trail, and compliance challenges: Current approaches to LLM security are unlikely to meet specific GxP requirements, because LLM outputs could become quality records subject to regulatory inspection. To close this gap, ZT-LLM-COMPASS proposes the Zero-Trust Compliance Architecture (ZCA). The four evolving planes apply fine-grained authorizations, retrieval governance, function calling constraints, and tamper-obvious audit packages. The key innovation is suggesting to treat all LLM interactions untrusted and verified over time. Essentially, this instantiates human-in-the-loop gates, which allow the model to suggest compliance while refusing to guarantee it without explicit approval. It maps controls to architectural sensors and actuators and specifies the artifacts needed to show that the controls are in place. Pharmaceutical organizations are evolving towards generative AI, and architectures mapping technical capabilities to regulatory requirements are needed. The ZT-LLM-COMPASS specification provides computer-helped design support for productivity, data integrity, traceability and auditability to help meet regulatory requirements and ensure patient safety.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

1. Anjaneyulu Muppalla, et al., "Artificial Intelligence in Regulatory Compliance: Transforming Pharmaceutical and Healthcare Documentation," International Journal of Drug Regulatory Affairs, 2025. Available: https://www.ijdra.com/index.php/journal/article/view/764/407
2. Kevin Young, et al., "AI-Enhanced Zero Trust Architecture for Enterprise Systems," ResearchGate, 2025. Available: https://www.researchgate.net/publication/397880920_AI-Enhanced_Zero_Trust_Architecture_for_Enterprise_Systems
3. Zihan Zhu, "Intelligent information management enables quality-by-design in pharmaceutical production," Scientific Reports, 2025. Available: https://www.nature.com/articles/s41598-025-27879-w
4. Martin Haimerl and Christoph Reich, "Risk-based evaluation of machine learning-based classification methods used for medical devices," BMC Medical Informatics and Decision Making, 2025. Available: https://link.springer.com/article/10.1186/s12911-025-02909-9
5. Ananya Goswami, "The Architecture of Generative AI Systems: What Enterprises Must Know," BigStep, 2025. https://bigsteptech.com/blog/generative-ai-enterprise-architecture-guide-2025
6. Kampanart Huanbutta, et al., "Artificial intelligence-driven pharmaceutical industry: A paradigm shift in drug discovery, formulation development, manufacturing, quality control, and post-market surveillance," European Journal of Pharmaceutical Sciences, 2024. Available: https://www.sciencedirect.com/science/article/pii/S0928098724002513
7. Intuition Labs, "Custom Pharma Software Design: A GxP Compliance Guide." Available: https://intuitionlabs.ai/articles/custom-pharmaceutical-software-design
8. Kavitha Palaniappan, et al., "Global Regulatory Frameworks for the Use of Artificial Intelligence (AI) in the Healthcare Services Sector," Healthcare (Basel), 2024. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC10930608/
9. David B Isaacks and Andrew A Borkowski, "Implementing Trustworthy AI in VA High Reliability Health Care Organizations," Fed Pract, 2024. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC11147434/
10. Jakob Mökander, et, "Challenges and best practices in corporate AI governance: Lessons from the biopharmaceutical industry," Front Comput Sci, 2022. Available: https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2022.1068361/full