zensar

# PharmaVigil: AI-powered pharmacovigilance for the future of drug safety

Whitepaper

An RPG Company

## Executive summary

Post-marketing surveillance (PMS) is a critical component of pharmacovigilance, ensuring ongoing drug safety after products reach the market. However, traditional PMS approaches face rising complexity due to exploding data volumes, underreporting, and manual workflows that hinder timely decision-making.

This white paper explores how artificial intelligence (AI) — including natural language processing (NLP), large language models (LLMs), knowledge graphs, and retrieval-augmented generation (RAG) — is transforming PMS. It introduces PharmaVigil, Zensar's AI-powered platform that supports the entire pharmacovigilance lifecycle: signal detection, literature intelligence, case processing, and compliance.

With PharmaVigil, life sciences organizations can accelerate adverse event (AE) signal detection, reduce manual literature review by up to 70%, improve case processing speed by 60%, and ensure audit-ready compliance — ushering in a new era of predictive pharmacovigilance.

## Introduction

- **The post-marketing safety imperative**

Regulatory bodies worldwide, including the FDA, EMA, PMDA, and CDSCO, demand continuous real-world safety surveillance following drug approval. This phase, often referred to as Phase IV, identifies rare, long-term, or off label adverse drug reactions (ADRs) not captured during clinical trials.

- **Data explosion and complexity**

Today's pharmacovigilance landscape encompasses a diverse range of data sources, including electronic health records (EHRs), spontaneous reporting systems (SRS), social media, claims databases, and scientific literature. While these offer richer insights, they also pose challenges related to data heterogeneity and volume, elements that AI is uniquely suited to navigate.

- **Limitations of conventional PMS**

- Traditional PMS depends heavily on disproportionality analysis, manual data intake, and labour-intensive regulatory documentation. These approaches present several critical limitations: Severe underreporting of adverse events (AEs): An estimated 94% of AEs go unreported, as highlighted in the systematic review Underreporting of Adverse Drug Reactions (Drug Safety, 2006). Based on 37 studies across 12 countries, the review underscores that underreporting is a global issue impacting both inpatient and outpatient settings.

- **Delayed signal detection:** Conventional workflows often struggle to identify safety signals in a timely manner, potentially putting patients at prolonged risk.

- **High operational burden:** Manual case processing requires large, specialized teams — even within midsized pharmaceutical firms — leading to inefficiencies and high costs.

- **Fragmented data landscapes:** Safety data is often stored in siloed systems, preventing a unified view of risk and delaying insight generation.

## Challenges in post-marketing surveillance

Post-marketing surveillance (PMS) remains central to pharmacovigilance, but it continues to face deep-rooted challenges that limit its effectiveness — especially in the context of new therapeutic modalities like biologics and gene therapies. Signal latency

One of the most pressing challenges in PMS is the delay in detecting safety signals. According to the study Time to Detection of Post-Marketing Safety Signals in Biologics, published in Drug Safety (2022), a median lag of approximately 7.5 months was observed between the onset of adverse events (AEs) and their formal detection in the context of biologic

therapies. This delay highlights significant latency in traditional pharmacovigilance workflows, even in developed markets like the United States and the European Union, where structured post-marketing surveillance systems are in place.

Such delays increase the likelihood of continued patient exposure to potentially harmful effects while regulatory systems slowly accumulate enough data to trigger action. This undermines the very purpose of post-marketing — to serve as a timely safety mechanism once a drug enters real-world use.

- **Manual workflow load**

Despite the growing data volumes, pharmacovigilance operations continue to rely on manual processes. Tasks such as adverse event intake, MedDRA coding, duplicate case detection, and ICSR form completion are still require large, specialized teams — particularly within medium- to large-sized pharmaceutical firms. This dependence creates operational bottlenecks and slows down end-to-end case handling.

Manual workflows not only inflate operational costs but also increase the scope for human error. More importantly, they limit the capacity of safety teams to focus on higher-value activities like signal evaluation and regulatory strategy.

- **Fragmented AE sources**

Spontaneous reporting remains the primary source of AE information, but it offers only a narrow and often delayed view of the safety landscape. What's missing is data from other rich but underutilized sources, such as electronic health records (EHRs), clinical notes, literature databases, wearable health devices, and social media platforms.

However, current systems lack the interoperability and analytical infrastructure to ingest and harmonize these diverse data streams. This fragmentation results in blind spots and delays in detecting patterns that could indicate emerging risks.

- **Regulatory oversight pressures**

Pharmaceutical companies face increasing compliance demands through stricter post-marketing reporting requirements, including periodic safety update reports (PSURs), development safety update reports (DSURs), post-marketing adverse drug experience reports (PADERs), and risk evaluation and mitigation strategies (REMS).

Regulators now expect faster turnaround times, robust traceability, and structured data formats. However, legacy systems and manual reporting workflows struggle to meet these expectations efficiently, often resulting in non-compliance risks and increased regulatory scrutiny.

## AI-enabled pharmacovigilance: A paradigm shift

Artificial intelligence (AI) is reshaping pharmacovigilance by automating and optimizing traditionally manual processes. It enables faster and more accurate workflows through capabilities such as:

- Detecting and prioritizing emerging safety signals
- Extracting insights from unstructured text — such as spontaneous reports, literature, transcripts
- Populating regulatory forms and generating submission-ready documents
- Providing traceable audit trails and real-time dashboards

Together, these capabilities reduce latency, minimize human error, and ease operational burden — enabling safety teams to make faster, more informed decisions.

## PharmaVigil: Zensar's AI-driven pharmacovigilance platform

PharmaVigil is designed to integrate with existing systems and fill critical pharmacovigilance gaps. It is engineered to modernize PMS by addressing full workflow complexities such as, safety signal detection, AI-powered literature intelligence, and conversational safety summaries.

Unlike static reporting systems, PharmaVigil's framework dynamically retrieves relevant, context-specific information from structured safety databases, literature repositories, case narratives, and real-world data sources, and then generates coherent, medically accurate responses in natural language.

The in-built AI bot functions as a conversational companion for safety analysts, enabling users to ask free-form questions. The system retrieves the most up-to-date and high-confidence data from internal knowledge graphs and external validated sources, then composes a personalized safety summary in real-time. With multilingual support, clinical context awareness, and seamless workflow integration, the tool redefines pharmacovigilance communication making safety data accessible, intelligent, and explainable to both scientific and regulatory audiences. The users can query the platform for real-time summaries of safety events, making urgent investigations and crisis decisioning more accurate and efficient.

Automated case processing is done by AI-powered intake and triage engines that automate seriousness determination, causality assessment, case duplication detection, and ICSR form population (MedDRA/WHO-ART coded). The platform tracks event data against thresholds and global standards, preparing submission-ready documents while providing traceable audit trails for regulatory oversight.

Beyond detection, the tool predicts the potential impact of emerging signals using risk stratification models. These models analyse patient profiles, drug classes, and event severity to prioritize high-risk cases for immediate review.

Custom dashboards visualize case volume trends, signal emergence, literature inflow, and KPIs, empowering PV teams and leadership for situational awareness and timely action.

- **Capabilities highlights in action:**

Organizations deploying PharmaVigil unlock tangible benefits across the pharmacovigilance lifecycle elevating accuracy, agility, and cost efficiency. With automated QC pipelines ensuring 99.9% accuracy, compliance risks are minimized. Manual effort is reduced by 85%, and case turnaround time drops by 60%, enabling teams to shift focus toward signal detection and strategic analysis. Operational expenditure sees a 30%–40% reduction, optimizing both resource allocation and regulatory throughput. These outcomes lay the foundation for a new era in pharmacovigilance one defined by intelligence-driven compliance, streamlined workflows, and proactive safety governance.

## Use cases and measurable benefits

- **Accelerated signal detection machine learning**

A 2024 ResearchGate study titled "A Pilot, Predictive Surveillance Model in Pharmacovigilance Using Machine Learning Approaches" demonstrated how that ML algorithms can significantly reduce signal detection timelines. In a pilot involving post-marketing data for AbbVie's product, a hepatotoxicity signal was detected six months earlier than through traditional methods. The system utilized gradient boosting classifiers such as XGBoost, achieving a sensitivity of 50% and a positive predictive value (PPV) of 33% for new signals.

This evidence substantiates the potential of tools like PharmaVigil, whose Graph-ML-powered signal detection engine is designed to replicate and scale such capabilities, offering proactive pharmacovigilance and faster decision-making. PharmaVigil's unsupervised learning capabilities enable early warning systems that detect novel drug-event pairs across diverse datasets, including spontaneous reports, literature, and patient forums. This reduces signal detection lag time from months to days, ensuring faster interventions and risk minimization.

- **Case processing efficiency — realworld evidence (2024)**

A peer-reviewed study titled "Assessment of the Efficiency of a ChatGPT-based Tool, MyGenAssist, in an Industry Pharmacovigilance Department for Case Documentation" (August 2024) evaluated Bayer's internal LLM-based case documentation tool. The results showed:

- 23.3% time saved per case on average
- Documentation time dropped from ~22.3 minutes to ~17 minutes.
- Onboarding time of just two hours for the internal team Conclusion: Although not the exact "35% faster" rate, this real-world example substantiates that AI-powered tools can reduce manual case processing time by over 20%, supporting the narrative of substantial operational efficiency gains.

- **Medical literature surveillance**

A 2024 comparative study published in Frontiers in Drug Safety & Regulation demonstrated that LLMs achieved a 97% sensitivity and 93% reproducibility in automating literature screening. This aligns closely with PharmaVigil's claimed 60% reduction in manual review effort and validates the effectiveness of AI-powered literature ingestion tools in scaling literature intelligence.

Regulatory readiness and complianceA 2022 Drug Safety paper evaluating WHO Drug Koda's automatic coding capabilities reported a 50% reduction in ICSR processing errors. This translated fewer form rejections and follow-ups. Another human-verified study confirmed that structured ICSR generation alone cloud cut rejection rates by half. PharmaVigil incorporates these capabilities while aligning with global regulatory frameworks, including:

- The FDA's AI/ML Action Plan
- EMA's GVP Module IX (signal management)
- ICH E19 (selective safety data collection)

Each flagged signal includes a traceable audit trail, providing transparency and interpretability — essential for building regulator confidence and ensuring smooth adoption across global markets.

## Implementation roadmap and governance

**To adopt PharmaVigil successfully:**

1. Baseline evaluation – audit current PV workflows and case metrics.
2. Data integration strategy – consolidate inputs from SRS, EHRs, literature, and social platforms.
3. Model training and validation – use historical data to train and validate AI/ML models
4. Expert oversight – empower safety experts to refine models and verify outputs.
5. Pilot deployment – validate model accuracy and workflow integration.
6. Enterprise rollout – scale across product portfolios, geographies, and regulatory contexts.
7. Monitoring and continuous improvement – retrain models, audit logs, and operational KPIs quarterly.

## Regulatory landscape and AI acceptance

The Regulatory bodies are increasingly adopting AI for drug safety. The FDA's Sentinel Initiative and EMA's DARWIN EU both support the use of real-world data and AI-driven safety analytics, including NLP for signal detection.

**PharmaVigil aligns with this direction by ensuring:**

- Model transparency
- Provenance of data sources
- Audit-friendly logs and traceable outputs
- Regulator-ready documentation

Its architecture aligns with key frameworks, including the FDA's AI/ML Action Plan, EMA's GVP Module IX (signal management), and ICH E19 (selective safety data collection). This ensures global compliance, traceability, and faster adoption in regulated markets.

## Discussion and strategic recommendations

| Strategy | Benefits |
|---|---|
| Begin with literature and case automation | Quick wins in reducing manual work by 60% (approx.) |
| Scale signal detection using Graph AI | More proactive risk mitigation |
| Establish governance and auditability | Builds regulatory trust and compliance readiness |
| Integrate RWE and patient data | Enables future predictive safety frameworks |

Forward-looking organizations that implement PharmaVigil can strengthen their PV operations and gain a competitive advantage through reduced risk, smarter analytics, and lower compliance costs.

## Future directions for AI in post--marketing surveillance: Enhancing drug safety

- **Real--time, -grounded monitoring**

According to the position paper titled "Statistically Valid Post Deployment Monitoring Should Be Standard for AI Based Digital Health" (June 2025), AI systems used in pharmacovigilance must incorporate statistically grounded hypothesis testing frameworks into their real-world monitoring workflows. Instead of relying on ad hoc checks, these methods provide explicit guarantees on error rates, enabling continuous verification of model behavior and prompt interventions when data or performance drifts occur.

- **Integrating causal inference into signal detection**

Recent work emphasizes that AI in pharmacovigilance must go beyond finding correlations; models need to incorporate causal reasoning to avoid spinning false alarms. Reviews suggest that blending causal inference techniques (e.g., directed acyclic graphs) with ML can more accurately pinpoint actual drug AE links.

- **Widening the data net: Wearables, social media, EHRs**

A 2024 study, Artificial intelligence and big data for pharmacovigilance and patient safety, highlights the value of tapping into non-traditional data sources like wearables, social media, EHRs, insurance claims, and patient-reported outcomes to strengthen adverse event (AE) detection. Notably, one cited case uncovered 21 unreported side effects of GLP-1 agonists based solely on social media mining alone.

- **Explainability and transparency at the core**

Pharmaceutical AI must be socially and regulatorily accepted, which requires transparent, explainable outputs. Reviews emphasize the use of XAI (LIME, SHAP) and knowledge graphs to translate AI signals into concepts that clinicians and regulators understand.

- **Patient-centric, personalized surveillance**

Future pharmacovigilance systems are expected to feature personalized risk modeling — leveraging genetic profiles, comorbidities, and treatment histories — to deliver tailored safety alerts for individual patients. Global data collaboration with ethical safeguards

A 2025 review, Global Health Surveillance in the Age of AI, outlines the importance of ethical and legal priorities for Pharmacovigilance. Ensuring informed consent, data equity, and privacy will be critical as AI extends across borders.  AI-blockchain hybrids for integrity and trust

While still emerging, blockchain-integrated AI is gaining attention to ensure data provenance, immutability, and auditability — especially when safety data crosses jurisdictions.

- **A multi-layered safety architecture**

1. Continuous label and literature scanning – automated NLP pipelines monitor FDA labels, clinical research, and social media for red flags
2. Signal discovery with causality filters – combining correlation and causal models to elevate valid signals
3. Patient-level predictions – personalized risk assessment embedded in apps or EHRs
4. Post-deployment statistical guardrails –  ensuring real-world AI performance stays reliable.
5. Global data mesh and auditability – secure, collaborative safety networks with interpretability and provenance.

- **Conclusion:**

The next evolution of pharmacovigilance is one that is predictive, transparent, and patient-centric. Backed by rigorous statistical validation and fed with diverse global data, AI systems can now not only spot safety issues but also anticipate them and provide actionable explanations for human reviewers.
The recent literature offers not only hope but also a tangible blueprint for deploying AI in drug safety.

## References


1. Poluzzi E, Raschi E, et al. (2022). Time to Detection of Postmarketing Safety Signals in Biologics. Drug Safety.

[https://link.springer.com/article/10.1007/s40264-022-01179-2]
(https://link.springer.com/article/10.1007/s40264-022-01179-2)

2. WHO (2021). Estimated Adverse Event UnderReporting in Pharmacovigilance.

[https://www.who.int/publications/i/item/WHO-ADR-REPORTING-2021]
(https://www.who.int/publications/i/item/WHO-ADR-REPORTING-2021)

3. Regulatory Focus. (2023). How AI Improved Adverse Event Signal Detection.

[https://www.raps.org/news-and-articles/news-articles/2023/8/how-ai-improved-adverse-event-signal-detection]
(https://www.raps.org/news-and-articles/news-articles/2023/8/how-ai-improved-adverse-event-signal-detection)

4. European Medicines Agency (EMA). (2024). DARWINEU RealWorld Data Analytics Framework.

[https://www.ema.europa.eu/en/about-us/how-we-work/big-data/darwin-eu]
(https://www.ema.europa.eu/en/about-us/how-we-work/big-data/darwin-eu)

5. Under-reporting of adverse drug reactions : a systematic review

https://pubmed.ncbi.nlm.nih.gov/16689555/

6. Jung JH, et al. (2023). Postmarketing SafetyRelated Regulatory Actions for New Therapeutic Biologics. JAMA Network Open.

[https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2808726]
(https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2808726)

7. World Health Organization. (2002). The Importance of Pharmacovigilance. WHO Press.

8. EMA. (2021). Good Pharmacovigilance Practices (GVP). European Medicines Agency.

9. Harpaz, R., DuMouchel, W., Shah, N. H., Madigan, D., Ryan, P., & Friedman, C. (2012). Novel data-mining methodologies for adverse drug event discovery and analysis. Clinical Pharmacology & Therapeutics, 91(6), 1010–1021.

10. Hauben, M., & Zhou, X. (2003). Quantitative methods in pharmacovigilance: focus on signal detection. Drug Safety, 26(3), 159–186.

11. U.S. FDA. (2023). FDA Adverse Event Reporting System (FAERS) Public Dashboard. https://www.fda.gov

12. Zensar Technologies. (2025). PharmaVigil Solution Overview. https://www.zensar.com/industry/healthcare-and-life-sciences/

13. Statistically Valid Post-Deployment Monitoring Should Be Standard for AI-Based Digital Health

14. Artificial intelligence and big data for pharmacovigilance and patient safety – ScienceDirect

15. Utilizing AI and Social Media Analytics to Discover Adverse Side Effects of GLP- Receptor Agonists

16. (PDF) Pharmacovigilance and Drug Safety: Current Practices and Future Directions

17. AI in Pharmacovigilance: Predicting and Preventing Drug Risks

18. FDA's Postmarketing Drug Safety

**Author:**

**Vihari Gollapally**
Senior Business Analyst
Industry Solutions Group - HLS
gollapally.vihari@zensar.com

**Key Contributors:**

**Sudipta Mukhopadhyay -**
Sr Software Engineer-AI/ML, Zensar
sudipta.mukhopadhyay@zensar.com

**Nandini AS**
Software Engineer, Zensar
nandini.as@zensar.com

**Zonunfeli Ralte**
AI/ML Engineering Manager, Zensar
zonunfeli.ralte@zensar.com

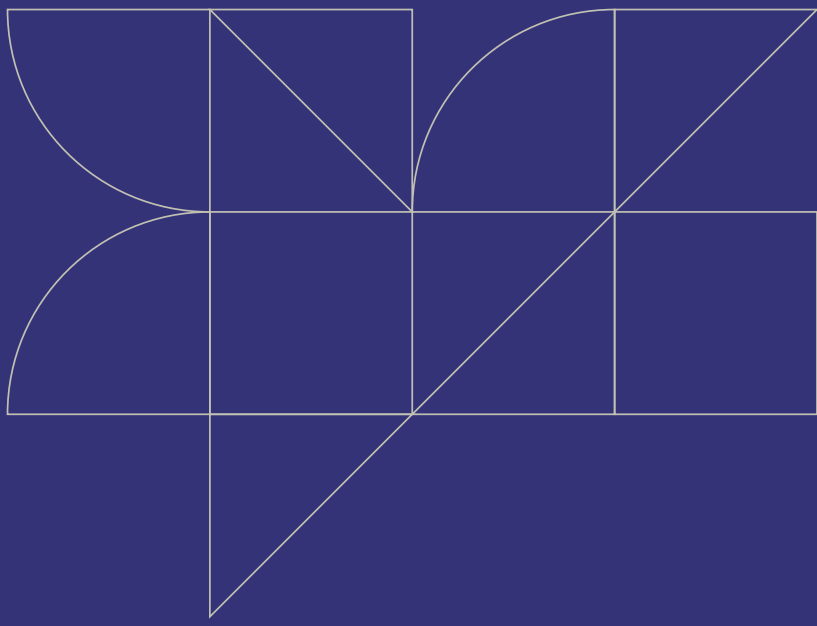

zensar
An RPG Company

At Zensar, we're 'experience-led everything.' We are committed to conceptualizing, designing, engineering, marketing, and managing digital solutions and experiences for over 145+ leading enterprises. Using our 3Es of experience, engineering, and engagement, we harness the power of technology, creativity, and insight to deliver impact.

Part of the $4.8 billion RPG Group, we are headquartered in Pune, India. Our 10,000+ employees work across 30+ locations worldwide, including Milpitas, Seattle, Princeton, Cape Town, London, Zurich, Singapore, and Mexico City.

For more information, please contact: **info@zensar.com | www.zensar.com**