CONTROLLED DIGITAL INTAKE

From social-media message to controlled action

An evidence-based pilot design for pharmacovigilance and product-quality intake using n8n and bounded AI

Reinhold Schilling

Correspondence and further publications: www.schilling-vigilance.com

24 September 2026

Citation: Schilling R. From social-media message to controlled action: an evidence-based pilot design for pharmacovigilance and product-quality intake using n8n and bounded AI. Schilling Vigilance White Paper Series, No. 5, Rev. 0. September 2026.

Abstract

Company-managed digital channels can contain suspected adverse reactions, exposure information, medication errors, quality defects and mixed safety–quality reports. The primary failure mode is not a wrong AI label in isolation but an interaction that is not captured, accepted by a qualified function or reconciled. This paper proposes an intended-use boundary, a source-to-disposition control chain and a risk-based silent-mode evaluation for an n8n-orchestrated intake workflow. It distinguishes routine channel surveillance from organised data collection under the now-applicable ICH E2D(R1), separates source posting from organisational awareness and ICSR reporting clocks, and treats urgent escalation and dual PV–Quality routing as independent controls. Release decisions are made separately for capture, routing, AI influence and customer-facing messages. The design is illustrative, not a validated deployment or a legal classification.

Keywords: pharmacovigilance; social media; ICH E2D(R1); product quality; n8n; AI; reconciliation; human oversight.

1. Why intake is a control problem

“The tablets looked strange and made me feel sick” is not a completed adverse-reaction report or a confirmed quality defect. It is information requiring prompt, coordinated assessment by pharmacovigilance (PV) and Quality. An intake system should preserve what the person actually said, where and when it appeared, what the organisation learned and did, and whether both functions accepted the handover. Neither a successful n8n execution nor an alert delivered to an inbox proves a regulated business outcome.

This design covers explicitly approved company-managed channels and messages accessible through approved platform interfaces. Channel-by-channel obligations, vendor arrangements, geographic coverage, moderation practices, deleted or edited posts, attachments and direct messages must be inventoried. “Every interaction” below means every in-scope interaction observable through the approved source and evidenced by the source-to-workflow reconciliation; platform data that cannot be accessed must be recorded as a limitation, not represented as complete capture. GVP Module VI addresses screening of internet or digital media under the marketing-authorisation holder’s management or responsibility. [1,2]

2. Regulatory position at publication

GVP Module VI revision 2 remains the EMA-listed adopted module as of 24 September 2026. ICH E2D(R1) took effect in the EU on 18 March 2026; its additional transition period ended on 18 September 2026. The EMA implementation strategy says the revised definitions and guidance are now applicable while Module VI revision is prepared. It specifically identifies social listening or digital listening as an example of an organised data collection system (ODCS), distinct from routine spontaneous-report surveillance. A pilot that goes beyond routine receipt and screening to planned analysis of user communications must determine and document whether it is an ODCS. For a new non-protocol ODCS, the strategy calls for documentation of objectives, sources, dataset and look-back/duration, review method, and management of safety observations, under QPPV oversight; ongoing ODCSs have a narrower minimum documentation set. Do not assume all company-page messages are solicited, or that all digital monitoring is routine. [1,3,4]

Capture at least three separate times when available: source post/publication time, first organisational awareness or technical ingestion time, and downstream PV receipt time. Qualified PV staff determine case validity, source classification, applicable day zero and reporting deadlines under current EU rules; an automation receipt timestamp must not silently reset a clock. The intake may flag incomplete reports for follow-up but must not wait for a questionnaire before forwarding potentially relevant information. [2,4]

The EU AI Act generally became applicable on 2 August 2026; some provisions have different dates, including revised high-risk deadlines of 2 December 2027 for Annex III and 2 August 2028 for certain Annex I systems under Regulation (EU) 2026/1744. These dates do not, by themselves, classify this intake tool as high-risk. The organisation should document its actual intended purpose, provider/deployer roles, AI-literacy measures and, if it deploys a conversational AI interface, applicable transparency duties before public use. [5,6]

3. Intended use and accountable roles

Proposed intended use: receive and preserve accessible in-scope channel interactions; detect and prioritise potential safety or quality signals; suggest a preliminary route; ask only approved, limited questions; and hand over the original content to qualified PV, Quality or Medical Information (MI) personnel. Excluded: determining ICSR validity, causality, seriousness, expectedness, reportability, final complaint or defect classification, treatment advice, and recall or market action.

The marketing-authorisation holder remains accountable for its PV system even when activities are contracted out. PV and the QPPV oversee safety procedures and relevant ODCS documentation; Quality owns defect triage and escalation; MI approves its responses; the channel owner verifies source access and counts; the technical owner maintains interfaces, credentials and recovery; and the quality unit oversees risk-based qualification, deviations and release. EMA inspection guidance stresses intended-purpose qualification/validation and assessment of vendor documentation rather than assuming that the vendor's product is already validated for the company's use. [3,7]

Control state	Accountable operational evidence
Captured	Source identifier, original text/media reference, version or edit history where available, source time and ingestion time.
Registered	Unique interaction ID, durable controlled source copy, access controls and linkage to workflow execution.
Triaged	Rule hits, AI proposal if used, language, urgency flag, exception state and versioned configuration.
Transferred	Recipients, transfer attempts, acknowledgements, downstream references and escalation timers.
Reconciled	PV and Quality dispositions where applicable; signed discrepancy resolution and closure record.

4. Workflow and decision boundaries

4.1 Capture before classification

A source connector ingests each observable in-scope event and creates a controlled record before calling a model. Preserve original wording and attachment references separately from derived summaries; retain edits, removals and corrections as traceable events where source access permits. Define deduplication by source ID and event version, retries with idempotent writes, and a manual capture route when the platform or connector is unavailable. If registration fails, alert an on-call owner from an independent monitoring path: a failed workflow cannot be expected to log itself reliably.

4.2 Priority, routes and AI

Use independent, multi-label PV, Quality and MI flags plus an orthogonal urgency flag; "emergency" must not replace the PV or Quality route. Configure conservative indicators for death, breathing difficulty, severe allergic reaction, hospitalisation, overdose, accidental paediatric exposure, wrong medicine or strength, contamination, falsification, tampering and possible multi-patient impact. These are screening indicators, not regulatory seriousness or defect classifications. Negation, quotations and hypothetical wording require testing so that "I was not hospitalised" does not automatically become a confirmed hospitalisation. Human reviewers adjudicate uncertain meaning.

Deterministic triggers are positive safety rails, not an exhaustive detection guarantee: a symptom can be described without a product name and a defect without obvious keywords. A bounded model may propose zero, one or multiple routes and a reason code using a fixed schema, but it cannot clear a hard-rule flag. Malformed, unavailable, contradictory, low-confidence or unsupported-language outputs are queued for human review. No "general enquiry" label may silently close an item in the pilot. Treat source text as

untrusted input: block prompt-injection instructions, restrict model tool access and test for leakage or alteration of sensitive content.

4.3 Transfer and follow-up

For any potential PV or complaint content, dispatch the available original promptly to the appropriate function; a follow-up exchange runs in parallel, never as a prerequisite. Use approved questions for product, occurrence, person affected, dates and recontact for safety matters, and product/strength, defect, batch, expiry, sample, use and health outcome for complaints. Avoid soliciting unnecessary identifying details in public comments; direct people to an approved private channel when appropriate. Two linked downstream records and acknowledgements are required for a combined route. Each pending route remains open with an owner, a risk-based acknowledgement timer and a tested backup path.

A customer-facing response is a separate release scope. Restrict it to approved acknowledgement, contact and emergency signposting templates; do not let a classifier invent medical advice. Verify whether the interface must disclose that the person is interacting with AI, and preserve human escalation for urgent contacts. [5]

5. Silent-mode protocol

Keep the existing manual process authoritative. Shadow-process the same in-scope messages without autonomous closing or customer replies, plus a pre-approved synthetic challenge set. PV and Quality reviewers create an independent reference decision, blinded to AI output where feasible, and adjudicate disagreements with documented reasoning. Label each record for potential PV, potential quality complaint, urgency, required recipients, supported language, source type, expected follow-up and disposition. Separate routine screening from any planned ODCS activity in the pilot protocol. [4]

Stratify the challenge set across direct and indirect symptoms, no product name, brand variants, negation, hypothetical statements, emojis and typos, multilingual content, missing reporter details, deleted/edited posts, attachments, combined PV–Quality cases, wrong-strength packs and connector failures. Include naturally occurring pilot traffic to estimate workload and case mix; synthetic challenges cannot by themselves demonstrate production prevalence or sensitivity.

Measure	Definition / release interpretation
Capture completeness	Matched observable source events / all observable in-scope source events; investigate every unmatched event and define source-count limitations.
PV / Quality sensitivity	True potential cases routed to required function / all reference potential cases in that stratum; report counts and uncertainty, not accuracy alone.
Combined-route completion	Both functions acknowledged / reference combined cases; review each miss individually.
Urgent acceptance latency	Time from first recorded organisational awareness to qualified urgent acceptance; report distribution and threshold breaches.
False-positive workload	Human-review volume by channel, language and shift; demonstrate staffing is adequate.
Exception recovery	Failed ingestion, handover or abandoned chats with documented ownership and final resolution.

Predefine sample size, strata, acceptance limits and escalation windows against channel volume and risk. A run with zero observed misses does not prove zero future misses; report the number tested and uncertainty, and keep human surveillance and reconciliation in place. Each mismatch records the original, expected and actual routes, possible patient/quality impact, root cause, corrective action, regression set and review of similar live interactions. A potential lost or delayed safety/quality item blocks the affected release scope until

its impact is assessed and effective controls are retested.

6. Failure control and reconciliation

Failure mode	Fallback and proof of resolution
Connector or registration failure	Independent source-to-record check; alert on-call owner; manually preserve message with original source time and reconcile backlog.
Model unavailable or unsafe output	Retain record and queue for human triage; record reason and configuration version.
One of two handovers fails	Keep combined interaction open; use approved manual route; require both downstream acknowledgements.
Unsupported language or incomplete chat	Preserve available content, obtain qualified language review where needed and route without waiting for completion.
Alert not accepted	Escalate from monitored queue to backup/on-call role; alert delivery alone is not acceptance.

Reconcile source event counts and IDs to durable interaction records, then reconcile each required route to acceptance and final disposition. Define frequency and after-hours coverage based on applicable reporting and defect risk, not a generic “each working day” rule for urgent items. Record open exceptions and age, owner, correction, evidence and sign-off. n8n’s logging and human-review facilities are implementation aids, not a substitute for controlled source archives, retained audit evidence, qualification or an operating procedure; execution-data retention and privacy settings need explicit verification in the chosen deployment. [8,9]

Quality-defect escalation needs its own risk-based SOP. EMA’s May 2026 procedure on suspected quality defects is addressed primarily to competent authorities; its authority-assessment timelines must not be copied into a company’s intake SLA. Company notification and market-action responsibilities must be mapped separately for the product and jurisdiction. [10,11]

7. Release and lifecycle governance

Release gate	Minimum evidence
Capture	In-scope source inventory, durable records, independent completeness check, failed-ingestion recovery.
Routing	Stratified reference comparison, urgent and dual-route tests, acknowledgement and backup evidence.
AI influence	Schema/fallback and prompt-injection tests, model/prompt version traceability, no suppression of deterministic flags.
Customer communication	Approved templates, human escalation, AI disclosure assessment and privacy-approved channel design.

No-go: any unresolved failure capable of losing, suppressing or materially delaying a potential safety or quality message in the proposed scope. A restricted release is possible only where exclusion by channel, language, function or message type is technically enforced, operationally staffed and monitored. The approved evidence pack includes intended use, source inventory, risk assessment, requirements-to-test traceability, qualified configuration, supplier and privacy/security assessment, ODCS determination where relevant, challenge data, protocol and results, deviation log, SOPs and training, reconciliation records, and a signed scope-specific release decision.

After release, monitor queue age, source-to-record gaps, false negatives from periodic human sampling of negative labels, changes in product terms and languages, model drift, prompt or vendor changes, and downstream acceptance. Changes to rules, prompts, models, interfaces or destinations require impact assessment, controlled testing, approval, baseline update and heightened post-change monitoring. Limit personal-data transfer to necessary fields; assess legal basis, vendor processing and international transfers, retention, access and deletion alongside recordkeeping duties. If a case proceeds to EudraVigilance, apply the EMA's 2025 Module VI Addendum II masking guidance in the regulated case-reporting process, not by indiscriminately stripping the intake source record. [1,12]

Conclusion

The defensible claim for this design is not that an AI model reliably identifies every case. It is that an organisation can demonstrate, within a declared and tested channel scope, what it received, when it became aware, how a potential safety or quality concern was prioritised, who accepted it, and how exceptions were resolved. Human authority, a functioning fallback and source-to-disposition reconciliation remain essential even when classification performance is strong.

Scope, declarations and citation

This is a proposed EU-focused control model for human medicinal products, not a description of a deployed or validated system, legal advice or a certification of n8n or any model. Product-specific and national requirements, source access, reporting clocks and operational service levels require sponsor approval. Citation: Schilling R. From social-media message to controlled action: an evidence-based pilot design for pharmacovigilance and product-quality intake using n8n and bounded AI. Schilling Vigilance White Paper Series, No. 5, Rev. 0. September 2026.

References

1. EMA. Good pharmacovigilance practices (GVP), overview and current modules. Accessed 24 Sep 2026. https://www.ema.europa.eu/en/human-regulatory-overview/post-authorisation/pharmacovigilance-post-authorisation/good-pharmacovigilance-practices-gvp

2. EMA. GVP Module VI: collection, management and submission of reports of suspected adverse reactions (Rev. 2). 2017. https://www.ema.europa.eu/en/human-regulatory-overview/post-authorisation/pharmacovigilance-post-authorisation/good-pharmacovigilance-practices-gvp

3. EMA. ICH E2D post-approval safety data management: effective EU version. 2026. https://www.ema.europa.eu/en/ich-e2d-post-approval-safety-data-management-scientific-guideline

4. EMA. EU implementation strategy of ICH E2D(R1). EMA/11141/2026. https://www.ema.europa.eu/en/documents/scientific-guideline/eu-implementation-strategy-ich-e2dr1-guideline-post-approval-safety-data-definitions-standards-management-reporting-individual-case-safety-reports_en.pdf

5. European Commission. AI Act: application timeline and transparency. Accessed 24 Sep 2026. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

6. European Union. Regulation (EU) 2026/1744, amending Regulation (EU) 2024/1689. https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng

7. EMA. Coordination of pharmacovigilance inspections: electronic systems and outsourced PV. https://www.ema.europa.eu/en/coordination-pharmacovigilance-inspections-0

8. n8n. Set up logging. Documentation. Accessed 24 Sep 2026. https://docs.n8n.io/deploy/host-n8n/keep-n8n-running/set-up-logging

9. n8n. Human-in-the-loop for tools; privacy and security. Documentation. https://docs.n8n.io/build/integrate-ai/ai-examples/human-in-the-loop-for-tools

10. European Commission. Management and classification of reports of suspected quality defects. May 2026. https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/management-classification-reports-suspected-quality-defects-medicinal-products-risk-based-decision-making_en.pdf

11. EMA. GMP/GDP questions and answers: reporting suspected quality defects. https://www.ema.europa.eu/en/human-regulatory-overview/research-development/compliance-research-development/good-manufacturing-practice/guidance-good-manufacturing-practice-good-distribution-practice-questions-answers

12. EMA. GVP Module VI Addendum II: masking of personal data in ICSRs submitted to EudraVigilance. 2025. https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-addendum-ii-masking-personal-data-individual-case-safety-reports-submitted-eudravigilance_en.pdf