NEXTVISE

PUBLIC WHITE PAPER / SEPTEMBER 2026

The Assurance-to-Runtime Gap

How regulated AI moves from requirements to verified decisions

Regulatory assurance. Domain verification.

Deterministic runtime enforcement.

THE OPERATING PRINCIPLE

An approved system is not enough. The consequential output must also be supported, controlled and reconstructable.

PURPOSE

Beyond governance documentation

Regulated institutions do not only need to show that an AI system has policies, owners and risk files. They need to know whether a consequential output is supported by the domain facts required for its intended use—and whether an unsafe or unsupported output can be intercepted before it becomes an action.

This paper explains the Nextvise operating model: regulatory assurance against applicable requirements, domain-grounded verification against controlled terminology and authoritative sources, and deterministic enforcement of approved conditions during use.

That distinction matters in healthcare and life sciences. A well-documented system can still misstate a diagnosis, confuse a medicine, omit a suspected adverse event, overstate a genomic association or create an uncontrolled GxP record. Explainability may describe an output; it does not establish that the output is clinically, scientifically or procedurally supported.

Nextvise follows the decision from obligation to use. It helps organisations prepare an AI management system, tests the supporting material, converts approved controls into operational conditions and preserves what happened. The paper is technology-neutral and does not disclose proprietary architecture.

CENTRAL THESIS

Nextvise does not stop at governance documentation. It verifies whether regulated AI outputs are supported by the domain information required for their intended use.

CONTENTS

Inside the paper

01	Why high-risk outputs fail
02	What Nextvise actually does
03	Regulatory assurance
04	Domain-grounded verification
05	Healthcare verification
06	Pharmacovigilance verification
07	GMP and draft Annex 22
08	Precision medicine and genomics
09	Human in the Loop
10	Deterministic runtime enforcement
11	An end-to-end safety case
12	Legal and auditor ecosystem
13	Decision records and deployment
14	Control reference matrix

01 / WHY HIGH-RISK OUTPUTS FAIL

A governed system can still produce an unsafe decision

The regulatory concern is not AI in the abstract. It is the use of AI where an unsupported output can alter care, conceal a safety signal, create a controlled record, deny an opportunity or mislead a professional reviewer.

External validation of a widely implemented sepsis model found poor discrimination and calibration.[13] A population-health algorithm underestimated the needs of Black patients because healthcare cost was used as a proxy for illness.[14] FDA recorded the use of AI agents to create drug specifications, procedures and production or control records during an inspection identifying CGMP violations.[15] The High Court of England and Wales examined legal submissions containing false or fictitious authorities.[16]

These are different failures, but they share an operational lesson: a policy, model card or general explanation does not prove that the specific output was fit for the specific decision.

> **THE GAP**
>
> Governance asks whether the organisation has a controlled system. Domain verification asks whether this output is supported. Runtime enforcement determines whether it may proceed.

Why regulation reaches beyond accuracy

- EU AI Act Articles 9–15 address risk management, data governance, documentation, records, transparency, human oversight, accuracy, robustness and cybersecurity.[1]
- ISO/IEC 42001 establishes a management system for governing AI across organisational processes.[2]
- Sector rules add the meaning of correctness: clinical context, pharmacovigilance reporting, GxP control, credit reasons or verified legal authority.

02 / THE NEXTVISE OPERATING MODEL

From requirement to a verified decision

Nextvise is infrastructure for defensible AI. It connects three disciplines that are usually separated: regulatory assurance, specialist verification and operational enforcement.

1	REQUIREMENT	Identify the exact law, standard, guidance, contract or approved internal control.
2	DOMAIN REFERENCE	Select the controlled terminology, validated dataset or authoritative source relevant to the intended use.
3	VERIFY	Test the material or output; identify contradiction, omission and unsupported claims.
4	DISPOSITION	Approve, Escalate or Block according to the approved control condition.
5	HUMAN IN THE LOOP	Route the case to a qualified person when judgement or consequence requires it.
6	DECISION RECORD	Preserve rule, source, version, reason, action, identity and time.

Deterministic does not mean that medicine, law or science becomes mechanically certain. It means the control outcome is reproducible against approved criteria. Where facts are insufficient, conflicting or high-consequence, the correct result is Escalate—not invented certainty.

WHAT NEXTVISE IS NOT

It is not a diagnostic system, a substitute for qualified judgement, a certification body or a claim that every model output can be made medically true.

03 / REGULATORY ASSURANCE

Test the management system against the rule that applies

The first layer establishes whether governance and control design are ready for scrutiny. Nextvise scopes the organisation, system, intended use, role and jurisdiction, then maps each relevant requirement to policies, system descriptions, risk files, validation material and operating records.

For ISO/IEC 42001 readiness, this includes context and scope (Clause 4), leadership (5), planning (6), support (7), operation (8), performance evaluation (9) and improvement (10).[2] For high-risk systems under the EU AI Act, assessment can extend to Articles 9–15 and the provider or deployer duties that apply.[1]

The assurance cycle

- Define the system boundary, intended use, accountable roles and applicable framework.
- Turn the requirement into a precise assessment criterion.
- Compare controlled documentation and operating material with that criterion.
- Issue a reasoned finding linked to the exact source and observed condition.
- Specify remediation, owner, acceptance test and review authority.
- Confirm closure before the control becomes part of the approved operating boundary.

Legal partners and consultants can review applicability and interpretation inside this process. Auditors receive a structured, read-only path from requirement to conclusion. Certification remains the decision of an accredited certification body.

THE OUTPUT

A finding states what applies, what was observed, why the gap matters, what must change and how closure will be verified.

04 / DOMAIN-GROUNDED VERIFICATION

A compliant process does not make an unsupported output correct

The second layer asks a different question: is the claim, classification or recommended action supported by the domain information required for this use? This is where Nextvise goes beyond generic compliance tooling.

Verification is configured around intended use. It compares entities, terminology, values, temporal relationships and claims with controlled vocabularies, validated customer sources and authoritative external references. The aim is to surface contradictions, missing support, improper extrapolation and conditions that require qualified review.

What is checked

- Identity: is the patient, product, medicine, event, variant, document or authority correctly resolved?
- Terminology: is the concept represented with the controlled vocabulary required by the workflow?
- Context: are population, indication, units, timing, provenance and intended purpose preserved?
- Support: does the source substantiate the claim at the required level?
- Consistency: does the output contradict a source, material fact or approved rule?
- Consequence: would an error affect care, safety reporting, product quality, rights or liability?

A model may assist with extraction or generation. The disposition is applied against approved criteria. Unsupported confidence does not become evidence because the language is fluent.

THREE OUTCOMES

Approve when defined checks pass. Escalate when judgement or more information is required. Block when a prohibited or unsafe condition is met.

05 / HEALTHCARE VERIFICATION

Check the clinical assertion before it shapes care

Healthcare language is dense with abbreviations, negation, temporal qualifiers, dosage units and context-specific meaning. A plausible answer can still confuse a medicine, transform a historical condition into an active diagnosis, omit a contraindication or overstate what a source supports.

Nextvise evaluates the output where it could influence diagnosis, treatment, triage, coding or patient communication. Controls depend on intended purpose and clinical risk; Nextvise does not make the final clinical judgement.

A healthcare verification profile can test

- Clinical entities, controlled terms, synonyms, abbreviations and negation.
- Medicine identity, dose, route, units and relevant patient context.
- Whether a diagnostic assertion is historical, suspected, excluded or confirmed.
- Whether the source supports the statement for the relevant population and use.
- Contradictions between generated output, source record and approved clinical rule.
- Whether uncertainty, missing information or potential harm requires qualified review.

RxNorm provides normalized names for clinical drugs and links across drug vocabularies; other controlled terminologies may be required by the customer's workflow.[11] Nextvise uses the approved references selected for the use case—it does not substitute statistical confidence for clinical support.

SAFETY BOUNDARY

Nextvise verifies support and control conditions. A qualified healthcare professional remains responsible for clinical judgement where the workflow requires it.

06 / PHARMACOVIGILANCE

Detect the possible safety case without inventing causality

Pharmacovigilance is not simply document classification. A workflow must recognise a possible suspected adverse reaction, preserve source facts and route the information so the responsible safety function can determine validity, seriousness, expectedness, follow-up and reporting.

EMA GVP Module VI covers collection, data management and submission of individual reports of suspected adverse reactions and describes report validation and handling requirements.[7] MedDRA supplies standardised medical terminology used for regulatory information on medical products.[10]

A Nextvise pharmacovigilance profile can

- Detect a possible medicinal product and adverse event in structured or unstructured material.
- Normalize the product, event and medical concept against approved terminology.
- Preserve patient, reporter, event, product, timing and source information available in the record.
- Surface missing elements, duplicates, contradictions and ambiguous temporal relationships.
- Apply approved seriousness or escalation criteria without claiming final causality.
- Route the case and supporting excerpts to the authorised safety professional.

The result is not an autonomous regulatory conclusion. It is a controlled disposition: sufficient for the next approved step, insufficient and requiring follow-up, or unsafe to proceed without review.

WHY THIS MATTERS

A fluent summary can omit the one fact that changes safety handling. Verification keeps source, terminology, missing information and reviewer action connected.

07 / GMP AND DRAFT ANNEX 22

Manufacturing control is a separate regulatory track

Pharmacovigilance, precision medicine and pharmaceutical manufacturing are related domains, but they are not interchangeable control regimes. Nextvise keeps them separate.

EU GMP Annex 11 applies to computerised systems used in GMP activities and addresses lifecycle risk management, validation, data, audit trails, security, incident management and business continuity.[8] The European Commission's draft Annex 22, published for consultation on 7 July 2025, proposes expectations for AI in medicinal-product manufacturing; it remains draft guidance, not final law.[9]

For an AI-enabled GMP use case, assurance can test

- Intended use, criticality, ownership and authorised operating boundary.
- Data provenance, suitability, lifecycle control and performance criteria.
- Validation strategy, acceptance criteria and fitness for intended use.
- Change control for model, data, prompt, rule, integration and infrastructure.
- Human oversight, exception handling and restrictions on automated action.
- Controlled records, electronic signatures and audit trails where 21 CFR Part 11 applies.[6]

FDA has separately discussed potential AI applications and regulatory considerations in drug manufacturing.[17] Nextvise uses such material to inform scoped assurance; it does not represent a draft or discussion paper as binding law.

CLEAN BOUNDARY

GVP governs pharmacovigilance activity. Annex 11 and draft Annex 22 address computerised systems and AI in GMP manufacturing. One checklist cannot stand in for the other.

08 / PRECISION MEDICINE AND GENOMICS

Verify the association, context and level of support

Precision-medicine outputs can combine variants, biomarkers, disease states, medicines and population evidence. The risk is not only a fabricated citation. It is also a real source applied to the wrong variant, phenotype, population, sample or clinical question.

ClinVar aggregates information about genomic variation and its relationship to human health, with submitted classifications and supporting material.[12] It is one authoritative reference that may support a verification profile; it is not a universal answer for every genomic decision.

A precision-medicine profile can test

- Variant, gene, biomarker and disease identifiers against the approved reference.
- Reference assembly, transcript, allele, sample and patient context where relevant.
- Whether asserted clinical significance and evidence level match the cited record.
- Whether the output overgeneralises across populations, indications or therapeutic settings.
- Whether sources disagree, are outdated or lack support for the claimed action.
- Whether the result must be escalated to a qualified specialist.

Nextvise does not promise deterministic biological truth. It provides a deterministic control decision against defined references and criteria. If the scientific record is uncertain or conflicting, Escalate is the defensible outcome.

PRECISION WITHOUT OVERCLAIM

The control can be reproducible even when the science contains uncertainty. Uncertainty must be preserved instead of smoothed away by generation.

09 / HUMAN IN THE LOOP

Qualified intervention, not a ceremonial approval

A human is not an effective control merely because a screen contains an approval button. EU AI Act Article 14 requires oversight proportionate to risk, autonomy and context, including the ability to understand limitations, avoid automation bias, interpret outputs, disregard or override an output and interrupt operation where appropriate.[1] Article 26 requires deployers to assign oversight to people with the necessary competence, training and authority.[1]

Nextvise routes a case with the information needed to act

- The original output and relevant source context.
- The exact criterion that triggered escalation.
- The controlled terminology or reference result.
- The contradiction, omission, uncertainty or prohibited condition detected.
- Permitted actions and the consequence of each action.
- Reviewer identity, decision, reason and timestamp.

Healthcare cases go to qualified healthcare professionals; pharmacovigilance cases to the authorised safety function; GxP exceptions to the appropriate quality or process owner; legal-source questions to qualified legal reviewers. The role follows the decision, not a generic queue.

Once the authorised human decides, the decision is recorded and the workflow proceeds according to it. No second approval loop is added after Human in the Loop.

OPERATIONAL TEST

The reviewer must be able to Approve, amend, Escalate or Block before the consequential action occurs—and have the competence and authority to do so.

10 / DETERMINISTIC RUNTIME ENFORCEMENT

Carry the approved boundary into live use

Pre-deployment assurance establishes readiness under reviewed conditions. Production introduces new inputs, model versions, agents, tools and operational pressure. The approved rule must remain actionable when the relevant condition appears.

Nextvise applies the approved control where the output or action matters. A deterministic control means that a defined condition produces a defined disposition. Probabilistic detection may identify the condition, but it cannot silently rewrite the approved rule.

The three dispositions

- Approve — required checks are satisfied and the output may continue within the approved boundary.
- Escalate — information is missing, conflicting, uncertain or requires qualified judgement.
- Block — a prohibited condition, unsupported high-consequence assertion or unauthorised action prevents continuation.

This is not a claim that every response legally requires the same gate. Control depth is proportionate to role, risk, autonomy, intended use and downstream effect.

The same model may be treated differently across workflows: an internal research draft may be permitted with disclosure, while a patient-facing dosage statement or autonomous change to a controlled manufacturing record may require escalation or blocking.

OPERATIONAL PRINCIPLE

Regulation defines the obligation. Assurance approves the control. Runtime enforcement makes the approved disposition unavoidable when the condition occurs.

11 / END-TO-END SAFETY CASE

A suspected adverse event in free text

An AI assistant summarises a customer message for a pharmacovigilance team. The message describes a medicine, reaction and timing, but the generated summary omits the reaction and presents the case as a routine product inquiry.

1 · Scope	Recognise the approved pharmacovigilance use case.	GVP and customer safety criteria apply.
2 · Ground	Resolve product and event; retain source excerpt and timing.	The possible safety signal stays visible.
3 · Verify	Compare summary with source and required case elements.	Material omission is detected.
4 · Enforce	Apply the approved rule for an omitted possible adverse event.	Escalate; routine continuation is not allowed.
5 · Review	Present source, trigger and missing element to safety professional.	Human determines follow-up and handling.
6 · Record	Store rule, references, versions, disposition, reviewer and time.	The decision can be reconstructed.

The model may produce the initial summary, but it does not control the disposition. Source facts, approved terminology and the safety rule determine whether the output can move forward.

THE DIFFERENCE

A compliance checklist can show that a procedure exists. Nextvise also tests the live output against the condition the procedure was meant to control.

12 / LEGAL AND AUDITOR ECOSYSTEM

Make specialist review usable without replacing it

Legal, audit and advisory teams do not need another isolated industry page. They need a coherent place inside assurance. Their role is to test applicability, interpretation, conformity and the reliability of the record supporting a conclusion.

Nextvise connects the exact requirement to the system fact, source, finding, remediation and operating result. Legal-source verification can surface missing or unsupported authorities before generated material is relied upon; the 2025 Ayinde and Al-Haroun proceedings show the consequences of fictitious citations and failed verification.[16]

What each participant sees

- The organisation owns the system, data, controls and remediation decisions.
- Legal partners assess legal applicability and advise on interpretation or exposure.
- Consultants support scoping, implementation and readiness.
- Lead auditors and audit teams review the structured trail and test conformity.
- Certification bodies independently decide certification where a scheme applies.
- Nextvise supplies controlled assurance and verification infrastructure; it does not issue the legal opinion or certification.

Auditor access is read-only and scoped. Customer information remains within the agreed boundary, with controlled references and exports rather than unrestricted operational access.

WHY LEGAL BELONGS IN ASSURANCE

The legal function establishes why a requirement applies and whether the record can withstand challenge. It is not a separate Nextvise product or vertical.

13 / DECISION RECORDS AND DEPLOYMENT

Defensibility without surrendering the data boundary

A log is not automatically a defensible record. The record must connect the governed condition to the information considered, the disposition and the accountable action.

A Nextvise decision record can preserve

- System, model, prompt, rule and reference versions.
- Input or source reference, with data minimisation and access control.
- Material output, detected condition and applicable requirement.
- Approve, Escalate or Block disposition and downstream status.
- Human reviewer identity, action, reason and timestamp where review occurred.
- Integrity, retention and export information required by the use case.

EU AI Act Articles 12, 19 and 26 establish record-keeping and log-retention duties in their respective high-risk provider and deployer scopes.[1] FDA-regulated electronic records may also engage 21 CFR Part 11 controls, including validation, protection, access, authority checks and time-stamped audit trails.[6]

Nextvise is model-, API- and agent-agnostic. Control can be deployed according to the customer's approved boundary, including cloud, sovereign, on-premises or isolated environments. Deployment location alone is not governance; identity, updates, references, logs and export paths must also be controlled.

DATA PRINCIPLE

The customer retains control of its data and operating boundary. Reviewers receive the scoped information needed for assurance, not unrestricted operational access.

14 / CONTROL REFERENCE MATRIX

Different questions require different sources

This matrix is a navigation aid, not legal advice. Applicability depends on role, jurisdiction, intended use and the specific decision.

AI management system	ISO 42001 Cl. 4–10; EU AI Act Arts. 9, 17	Are ownership, risk treatment, operation, evaluation and improvement established?
High-risk operation	EU AI Act Arts. 10–15, 19, 20, 26	Can data, output, oversight, logs and corrective action support intended use?
Healthcare output	Approved clinical sources and controlled terminology	Is the assertion correctly identified, contextualised and supported before use?
Pharmacovigilance	EMA GVP Module VI; MedDRA	Is a possible adverse reaction preserved, normalised and routed with required facts?
GMP / manufacturing	Annex 11; draft Annex 22; Part 11 where applicable	Are intended use, validation, change, records and human control approved?
Precision medicine	Approved genomic and biomarker sources; scoped ClinVar use	Does the claim match the record, context and evidence level?
Legal / professional	Applicable authority and controlled source material	Is the cited authority real, relevant and accurately represented?

The common discipline is traceability. The domain standard is not generic. Nextvise keeps the requirement, specialist reference, output test and final disposition in one defensible chain.

CONCLUSION

The output—not only the system—must be defensible

High-risk AI cannot be made safe by documentation alone. Governance establishes ownership, risk treatment and control. The organisation also needs to know whether the live output is supported in the domain where it will be used.

The Nextvise position is specific:

- Regulatory assurance tests the organisation and system against the applicable requirement.
- Domain-grounded verification checks critical outputs against approved terminology, sources and context.
- Deterministic runtime enforcement applies the approved disposition: Approve, Escalate or Block.
- Human in the Loop places ambiguous or consequential cases with the qualified professional.
- A controlled record connects requirement, source, output, disposition and accountable action.

This is particularly important in healthcare and life sciences, where an omitted event, wrong medicine, unsupported diagnosis, uncontrolled manufacturing record or overstated genomic association can create harm before a periodic audit detects the pattern.

Nextvise does not replace the clinician, safety professional, quality owner, lawyer or auditor. It gives each a more reliable point of intervention and a record showing why the decision was allowed to proceed.

> **NEXTVISE**
>
> Regulatory assurance plus domain-grounded verification plus deterministic runtime enforcement—infrastructure for defensible AI.

Important notice

This paper is general information, not legal advice, medical advice, an audit opinion or certification. Organisations must validate applicability and configure controls for their own system, intended use, jurisdiction and qualified professional responsibilities.

REFERENCES

Primary and authoritative sources

URLs accessed 30 September 2026. ISO/IEC 42001 is copyrighted; only clause-level references and paraphrases are provided.

[1] European Union, Regulation (EU) 2024/1689 (AI Act), especially Articles 9–20 and 26. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[2] ISO, ISO/IEC 42001:2023, Artificial intelligence—Management system. https://www.iso.org/standard/81230.html

[3] NIST, Artificial Intelligence Risk Management Framework 1.0, NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

[4] European Union, Regulation (EU) 2016/679 (GDPR), especially Articles 22, 25, 30, 32 and 35. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[5] eCFR, 12 CFR §1002.9 and official interpretations. https://www.ecfr.gov/current/title-12/chapter-X/part-1002/subpart-A/section-1002.9

[6] eCFR, 21 CFR Part 11, especially §§11.10, 11.50, 11.70, 11.100 and 11.200. https://www.ecfr.gov/current/title-21/chapter-I/subchapter-A/part-11

[7] EMA, GVP Module VI: Collection, management and submission of reports of suspected adverse reactions, Rev. 2. https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-collection-management-submission-reports-suspected-adverse-reactions-medicinal-products-rev-2_en.pdf

[8] European Commission, EudraLex Volume 4, EU GMP Annex 11: Computerised Systems. https://health.ec.europa.eu/medicinal-products/eudralex/eudralex-volume-4_en

[9] European Commission, Draft EU GMP Annex 22: Artificial Intelligence, consultation published 7 July 2025. https://health.ec.europa.eu/consultations/public-consultation-new-eudralex-volume-4-annex-22-artificial-intelligence_en

[10] MedDRA MSSO, Medical Dictionary for Regulatory Activities. https://www.meddra.org/

[11] U.S. National Library of Medicine, RxNorm. https://www.nlm.nih.gov/research/umls/rxnorm/index.html

[12] National Center for Biotechnology Information, ClinVar. https://www.ncbi.nlm.nih.gov/clinvar/

[13] Wong et al., External Validation of a Widely Implemented Proprietary Sepsis Prediction Model, JAMA Internal Medicine 181(8), 2021. https://doi.org/10.1001/jamainternmed.2021.2626

[14] Obermeyer et al., Dissecting racial bias in an algorithm used to manage the health of populations, Science 366(6464), 2019. https://doi.org/10.1126/science.aax2342

[15] U.S. FDA, Warning Letter: Purolea Cosmetics Lab, CMS 722591, 2 April 2026. https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/purolea-cosmetics-lab-722591-04022026

[16] High Court of Justice, Ayinde v Haringey and Al-Haroun v Qatar National Bank, [2025] EWHC 1383 (Admin). https://www.judiciary.uk/judgments/ayinde-v-london-borough-of-haringey-and-al-haroun-v-qatar-national-bank/

[17] U.S. FDA, Artificial Intelligence in Drug Manufacturing, CDER Discussion Paper, 2023. https://www.fda.gov/media/165743/download

[18] U.S. FDA, Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations, draft guidance, 2025. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing

[19] OWASP GenAI Security Project, Agentic AI—Threats and Mitigations. https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/