

SEC-049A - GAISSF Pharmaceutical Sector Executive Primer

Version: 1.1

Status: Controlled pre-release

Publisher: ODA3 Institute

Why this package exists

Pharmaceutical AI risk is not only a model-performance problem. It intersects with patient safety, product quality, clinical integrity, pharmacovigilance, regulated records, validated-state maintenance, cybersecurity and supplier dependence. SEC-049 provides an implementation layer between the authoritative GAISSF controls and pharmaceutical operating evidence.

What SEC-049 does not do

It does not determine GxP applicability, provide legal advice, replace a quality system, establish regulator approval, or guarantee safety, security or compliance.

Four implementation tiers

Tier 1: Administrative or low-impact support.

Tier 2: Controlled operational support.

Tier 3: GxP-significant or regulated decision support.

Tier 4: Safety-critical, quality-critical or high-autonomy regulated use.

The workbook score is indicative only. Qualitative GxP, patient, product, regulated-record and autonomy triggers can require escalation.

Five-step getting-started workflow

Inventory all AI systems, embedded components and external services.

Determine intended use, GxP status, regulated-record impact and physical/OT interaction.

Assign an initial criticality tier and obtain accountable approval.

Assess applicable GAISSF controls, document exceptions and link evidence.

Prioritise Tier 3/4 gaps, monitor changes and integrate incidents with deviation and CAPA.

First 90 days

During the first 30 days, identify unapproved regulated use, assign owners, classify high-risk systems and establish escalation. During days 31-90, complete GxP and criticality determinations, supplier reviews, control assessments, output-review rules, evidence capture and priority testing.

Notably Absent

NA-001. No reliable public evidence was identified that autonomous AI compromise of pharmaceutical manufacturing is widespread. Treat as a plausible high-impact scenario, not a prevalence claim. Confidence: Moderate; public reporting is incomplete.

NA-002. No claim is made that malicious manipulation of a pharmaceutical AI system has directly caused confirmed patient harm at scale. Do not infer occurrence from threat plausibility. Confidence: Moderate; confidential incidents may not be public.

NA-003. No evidence supports routine regulator acceptance of fully autonomous regulated decisions without accountable human and organizational controls. Default to explicit decision rights, qualified oversight and traceability. Confidence: High as a guide boundary; jurisdiction-specific review remains required.

NA-004. No universal global regulatory classification or validation method for pharmaceutical AI is assumed. Determine requirements by jurisdiction, intended use and lifecycle stage. Confidence: High.

NA-005. Public incident datasets do not provide complete coverage of AI failures in pharmaceutical operations. Frequency estimates are not supplied. Confidence: High.

NA-006. Conventional cybersecurity controls alone are not shown to be sufficient for GxP-relevant AI systems. Integrate quality, validation, data integrity, human oversight and regulated escalation. Confidence: High as an implementation principle.

NA-007. Model accuracy alone is not treated as evidence of clinical, safety, quality or regulatory fitness. Assess intended use, data, robustness, security, human factors, traceability and lifecycle control. Confidence: High.

NA-008. The guide does not establish that every listed threat has occurred. Threat entries are explicitly classified as scenarios unless confirmed evidence is cited. Confidence: High.

Release status

The package remains controlled pre-release. Qualified pharmaceutical quality/GxP, pharmacovigilance, clinical, manufacturing/laboratory, privacy, legal and jurisdiction-specific reviews remain required before public release.