HERMS WHITE PAPER | NOVEMBER 2026

Governed AI in Regulatory Affairs

A Framework for Human-Accountable Intelligence

HERMS Consultancy DMCC

Dubai, United Arab Emirates | hermsglobal.com

HERMS WHITE PAPER | WP5 | NOVEMBER 2026

Governed AI in Regulatory Affairs

A Framework for Human-Accountable Intelligence

Regulators in the United States, Europe, Saudi Arabia and Abu Dhabi now agree on one point: AI can support regulatory work, but people remain accountable for every decision. This paper turns that consensus into a practical framework for regulatory affairs teams in the GCC.

HERMS Consultancy DMCC

Dubai, United Arab Emirates | info@hermsglobal.com | www.hermsglobal.com

Contents

Executive Summary

1. Why Regulatory Affairs Needs Governed AI

2. The Regulatory Landscape

3. Core Concepts

4. The HERMS Human-Accountable Intelligence Framework

5. Risk Tiers for Regulatory Use Cases

6. Human Decision Rights

7. Evidence, Traceability and Audit

8. Generative AI and Large Language Models

9. Implementation Roadmap

10. Pitfalls and Readiness Checklist

11. How HERMS Helps

References

Executive Summary

AI is already part of regulatory work. FDA's drug centre reports experience with more than 500 submissions with AI components from 2016 to 2023.5 In January 2026, FDA and EMA published ten joint principles for good AI practice across the medicines lifecycle, the first of which is that AI should be human-centric by design.1

10 joint FDA–EMA principles for good AI practice, January 20261	**7 steps** in FDA's draft risk-based credibility framework3	**500+** submissions with AI components reviewed by FDA's drug centre, 2016–20235	**Oct 2025** Abu Dhabi Responsible AI Standard takes effect15

The direction is consistent across jurisdictions. Saudi Arabia's AI ethics principles require human oversight across the AI lifecycle and identifiable, accountable people.12 Abu Dhabi requires human-in-the-loop or human-on-the-loop oversight, override mechanisms and audit trails.15 EMA states that the applicant or marketing authorisation holder is responsible for ensuring that all algorithms, models and datasets are fit for purpose.7

Six commitments of governed AI

Define the context of use for every AI application before it is used.1,3

Set oversight in proportion to risk, based on how much the output influences a decision and what happens if it is wrong.3,4

Ground every output in sources that a reviewer can check.1,15

Assign a named human owner to every decision the AI supports.12,15

Log inputs, outputs, versions and human interventions for audit.15

Monitor and re-evaluate performance throughout the lifecycle.1,15

HERMS view

In regulatory affairs, the question is not whether AI is accurate enough to decide. It is whether the organisation can show who decided, on what evidence, and why.

1. Why Regulatory Affairs Needs Governed AI

Regulatory affairs is document-heavy, source-dependent and deadline-driven, which makes it well suited to AI support. It is also a function where an error can affect patient safety, product quality or a regulator's trust in the company.

Where AI adds value

Activity	Typical AI support	Human role
Regulatory intelligence	Monitoring authority websites, summarising new guidance, flagging changes	Confirm relevance and impact; decide action
Dossier preparation	First drafts, cross-checking modules, gap analysis against requirements	Review, correct and approve content
Labelling	Comparing SmPC, PIL and reference labels; tracking differences	Decide wording; approve local labels
Queries and deficiency letters	Locating prior responses and source data	Write and sign the response
Pharmacovigilance support	Case intake triage, literature screening	Medical assessment and causality
Tenders and pricing	Extracting requirements, checking document completeness	Commercial and compliance decisions

Where AI creates risk

Inaccurate output: EMA notes that large language models can produce variable results and irrelevant or inaccurate responses.8

Data leakage: the same EMA principles highlight potential data security risks and the need for safe input of data.8

Unclear accountability: if no person owns an AI-supported conclusion, nobody can defend it to a regulator.12

Silent drift: model performance can change over time, which is why the joint principles call for scheduled monitoring and re-evaluation.1

2. The Regulatory Landscape

There is no single AI rule for regulatory affairs. The framework in this paper draws on the following sources.

Source	Status	Relevance
FDA–EMA Guiding Principles of Good AI Practice1	Published January 2026; voluntary	Ten principles from human-centric design to lifecycle management; basis for future guidance2
FDA draft guidance on AI for regulatory decision-making3	Draft, January 2025	Seven-step, risk-based credibility framework for a defined context of use
EMA reflection paper on AI in the medicinal product lifecycle6	Adopted September 2024	Considerations for safe use of AI from discovery to post-authorisation
EMA/HMA guiding principles on large language models8	Published September 2024; updated regularly	Safe input, critical review of outputs, continuous learning, escalation
EU GMP Annex 22: Artificial Intelligence9	Consultation draft, July 2025	Draft excludes generative AI from critical GMP use; scope under review10
EU AI Act11	Phased application; high-risk rules from December 2027 and August 2028	Transparency rules apply from August 2026
SDAIA AI Ethics Principles and GenAI Guidelines for Government12,13	In force, Saudi Arabia	Human oversight; accountability; no GenAI in critical government decisions
DoH Abu Dhabi AI Policy and Responsible AI Standard14,15	Policy 2018; Standard October 2025	Risk tiers, human oversight, audit trails and documentation
ISO/IEC 4200116 and NIST AI RMF17	Published 2023	Management system and risk framework: govern, map, measure, manage

GCC implication

Companies operating in the Gulf should expect regulators to apply national AI ethics and health-sector standards alongside FDA and EMA expectations. A single governance model aligned with all of them is more efficient than separate controls per market.

3. Core Concepts

Context of use

FDA defines context of use as the specific role and scope of the AI model used to address a question of interest.3 The joint principles also require a well-defined context of use for every AI technology.1 In practice, this means one written statement per use case: what question the AI addresses, what it produces, and how the output will be used.

Model risk

FDA's draft guidance assesses model risk using two factors: model influence, meaning how much the output drives the decision, and decision consequence, meaning the impact of an incorrect decision.4 The higher the risk, the more validation, oversight and documentation are needed.3

Credibility

Credibility is trust, established through evidence, in an AI model's performance for a particular context of use.3 FDA's seven steps are: define the question, define the context of use, assess model risk, plan the credibility assessment, execute the plan, document results and deviations, and decide whether the model is adequate.3

Human in the loop and on the loop

Abu Dhabi's Responsible AI Standard requires human-in-the-loop or human-on-the-loop mechanisms, clear escalation protocols by risk level, and override mechanisms.15 Human in the loop means a person approves each output before use; human on the loop means a person supervises and can intervene.

4. The HERMS Human-Accountable Intelligence Framework

The framework below is a HERMS synthesis of the sources in Section 2. It organises AI governance in regulatory affairs into six pillars.

Pillar	What it requires	Anchored in
1. Purpose	A written context of use and question of interest for every AI application	FDA draft guidance3; joint principle 41
2. Proportion	Risk tier based on model influence and decision consequence; controls scale with tier	FDA draft guidance4; DoH risk classes15; joint principle 21
3. Provenance	Every output linked to identifiable, current sources; data lineage documented	Joint principle 61; DoH documentation rules15
4. People	Named owner for each AI system and each decision; defined decision rights; trained users	SDAIA accountability12; DoH roles15; joint principle 51
5. Proof	Logs of inputs, outputs, versions, reviews and overrides; audit-ready records	DoH traceability15; ISO/IEC 4200116
6. Performance	Pre-use validation, monitoring, drift checks and periodic re-evaluation	Joint principles 8 and 91; NIST AI RMF17

The accountability test

For any AI-supported regulatory output, the organisation should be able to answer within minutes: what was the AI asked, what sources did it use, who reviewed it, what did they change, and who approved it.

5. Risk Tiers for Regulatory Use Cases

The tiers below are a HERMS planning model. They apply the FDA logic of model influence and decision consequence4 and the Abu Dhabi principle that controls rise with risk.15

Tier	Examples	Minimum controls
Tier 1: Low	Internal search, meeting notes, first-pass translation for internal reading	Approved tool list; no confidential data in unapproved tools; user responsible for checking

Tier	Examples	Minimum controls
Tier 2: Moderate	Regulatory intelligence summaries, first drafts of non-critical sections, document completeness checks	Source citations required; qualified reviewer approves before use; sample audits
Tier 3: High	Content supporting safety, efficacy or quality claims; labelling text; responses to authority questions; PV case triage	Validated workflow; human approval of every output; full logging; documented credibility evidence; periodic re-evaluation
Tier 4: Not permitted	AI as sole decision-maker on submission, labelling approval, safety causality or batch-related decisions	Prohibited; AI may support but never replace the accountable person

Tier 4 reflects the wider regulatory position. Saudi Arabia's government GenAI guidelines bar generative AI from critical decisions affecting individuals and require human review of results.13 The draft EU GMP Annex 22 states that generative AI and large language models should not be used in critical GMP applications.9

6. Human Decision Rights

Governance fails when it is unclear who decides. Each AI-supported process should define four roles.

Role	Responsibility
System owner	Accountable for the AI system, its context of use, risk tier, validation and monitoring
Qualified reviewer	Checks each output against sources; corrects or rejects it; records the review
Decision owner	Approves the regulatory outcome, such as a submission, label or response, and signs it
Governance body	Approves new use cases and tier changes; reviews incidents, overrides and performance

Named people: SDAIA requires designers, developers and implementers to be identifiable and to assume responsibility.12

Defined monitoring duties: Abu Dhabi requires named individuals or teams to monitor outputs, review flagged cases and overrides, and report concerns.15

Final authority stays human: the Abu Dhabi standard states that clinicians should retain ultimate decision-making authority in AI-assisted healthcare scenarios.15 The same logic applies to regulatory sign-off.

Company responsibility: EMA places responsibility for the fitness of models and data on the applicant or marketing authorisation holder.7

7. Evidence, Traceability and Audit

A regulator or auditor may ask how an AI-supported conclusion was reached. The record should answer that question without relying on memory.

What to record

Sources: documents, versions and dates used to generate the output.1

Inputs and outputs: prompts or queries, outputs and the model or system version.15

Human actions: reviews, edits, overrides and approvals, each linked to a responsible person.15

Validation: test results, acceptance criteria and monitoring results.15

Data lineage: data provenance, processing steps and analytical decisions, documented in a traceable way in line with GxP.1

Retention and access

The Abu Dhabi standard requires version-controlled documentation available for internal and external audits, and logging of all access to AI systems.15 Regulatory teams should align AI records with existing document retention and quality-system rules.

Source-grounded by design

A regulatory AI output without a citation is an opinion. A cited output is evidence that a reviewer can verify. Governed systems should make citations mandatory for Tier 2 and Tier 3 use.

8. Generative AI and Large Language Models

Generative AI is the fastest-growing category in regulatory work and the one regulators treat most cautiously.

What regulators have said

EMA and HMA: four principles for staff: ensure safe input of data, apply critical thinking and cross-check outputs, keep learning, and know whom to consult when concerns arise. Organisations should define governance, permitted use cases, training and risk monitoring.8

EU GMP Annex 22 draft: generative AI and LLMs should not be used in critical GMP applications; in non-critical use, qualified people must confirm outputs are suitable.9 EMA is reviewing this boundary after industry feedback.10

SDAIA: government entities must not use generative AI in critical decisions affecting individuals, and all results should be human-reviewed.13

Upcoming EU guidance: the EU network workplan foresees guidance on AI in clinical development and pharmacovigilance.18

Practical controls

Use only approved, enterprise-controlled tools for confidential or product data.

Restrict generation to approved source libraries where possible, with citations.

Require a qualified reviewer to confirm every statement used externally.

Prohibit unreviewed AI text in submissions, labels and authority correspondence.

Test outputs regularly against known answers and record the results.

9. Implementation Roadmap

The roadmap is a HERMS planning model for a regulatory affairs function adopting governed AI.

Phase	Timing	Actions
Foundation	Days 0–90	Inventory current AI use; write an AI policy for regulatory affairs; approve tools; assign system owners; train users on safe input and review
Control	Days 90–180	Write context-of-use statements; assign risk tiers; set review and logging rules; define reviewer and decision-owner roles
Assurance	Days 180–365	Validate Tier 3 workflows; set acceptance criteria; start monitoring and sample audits; report to the governance body
Mature	Year 2 onwards	Align with ISO/IEC 42001; periodic re-evaluation; update for new FDA, EMA and GCC guidance

Measures of success

Share of AI-supported outputs with complete source citations.

Share of Tier 3 outputs with a recorded human review and approval.

Reviewer correction rate and time saved per task.

Number of audit findings related to AI use.

10. Pitfalls and Readiness Checklist

Ten pitfalls

Using AI without a written context of use.3

Applying the same controls to every use case regardless of risk.3

Entering confidential dossier or patient data into unapproved tools.8

Accepting outputs without checking sources.8

No named owner for the AI system or the decision.12

No log of who reviewed, edited or approved an output.15

Using generative AI for critical decisions.9,13

Validating once and never monitoring again.1

Training users on the tool but not on its limitations.8

Treating vendor assurances as a substitute for company responsibility.7

Readiness checklist

Area	Ready when
Policy	AI policy for regulatory affairs approved; permitted and prohibited uses listed
Inventory	Every AI use case recorded with owner, context of use and risk tier
People	Reviewers and decision owners named and trained
Evidence	Citations, logs and approvals captured for Tier 2 and Tier 3 outputs
Assurance	Validation, monitoring and re-evaluation schedule in place
Alignment	Controls mapped to FDA, EMA, SDAIA and DoH Abu Dhabi expectations

11. How HERMS Helps

HERMS Consultancy DMCC helps life sciences companies and healthcare organisations adopt AI in regulatory work with clear human accountability.

Governance design: AI policy, risk tiers, decision rights and governance bodies for regulatory affairs.

Use-case assessment: context-of-use statements, risk assessment and control design for each application.

Assurance: validation plans, acceptance criteria, monitoring and audit readiness.

Source-grounded regulatory intelligence: structured, cited regulatory content for GCC markets with human review.

Training: safe use, critical review and accountability for regulatory teams.

Request a governed AI readiness review

HERMS will map your current AI use against the framework in this paper and deliver a prioritised control plan.

info@hermsglobal.com | +971 55 600 8944 | www.hermsglobal.com

References

1. FDA and EMA. Guiding Principles of Good AI Practice in Drug Development, January 2026. https://www.fda.gov/media/189581/download

2. EMA. EMA and FDA set common principles for AI in medicine development, 14 January 2026. https://www.ema.europa.eu/en/news/ema-fda-set-common-principles-ai-medicine-development-0

3. FDA. Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products (draft guidance), January 2025. https://www.fda.gov/regulatory-information/search-fda-guidan…

4. DLA Piper. Key takeaways from FDA's draft guidance on use of AI in drug and biological life cycle, January 2025. https://www.dlapiper.com/en-at/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai

5. FDA CDER. Artificial Intelligence for Drug Development. https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development

6. EMA. Reflection paper on the use of Artificial Intelligence in the medicinal product lifecycle (EMA/CHMP/CVMP/83833/2023), September 2024. https://www.ema.europa.eu/en/scientific-guidelines/use-artificial-intelligence-ai-medicinal-product-lifecycle

7. EMA. Delivering on stakeholders' requirements for AI (presentation). https://www.ema.europa.eu/en/documents/presentation/presentation-delivering-stakeholders-requirements-ai-l-pinheiro-ema_en.pdf

8. EMA. Harnessing AI in medicines regulation: use of large language models, September 2024. https://www.ema.europa.eu/en/news/harnessing-ai-medicines-regulation-use-large-language-models-llms

9. European Commission. EudraLex Volume 4, Annex 22: Artificial Intelligence (consultation draft), July 2025. https://health.ec.europa.eu/document/download/5f38a92d-bb8e-…

10. EMA. GMP multistakeholder workshop on AI guidance development (Annex 22), 2026. https://www.ema.europa.eu/en/events/good-manufacturing-pract…

11. European Commission AI Act Service Desk. Timeline for the implementation of the EU AI Act. https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act

12. SDAIA. AI Ethics Principles. https://sdaia.gov.sa/en/SDAIA/about/Documents/ai-principles.pdf

13. SDAIA. Generative Artificial Intelligence Guidelines for Government. https://sdaia.gov.sa/en/SDAIA/about/Files/GenAIGuidelinesForGovernmentENCompressed.pdf

14. Department of Health Abu Dhabi. DoH becomes region's first entity to develop an AI policy in the healthcare sector, May 2018. https://www.doh.gov.ae/en/news/the-department-of-health-becomes-region

15. Department of Health Abu Dhabi. Responsible Artificial Intelligence (AI) Standard, DoH/ST/DDGO/RAI/V1/2025, October 2025. https://www.doh.gov.ae/-/media/B84FF5DD027846699C51F36C439A580F.ashx

16. ISO. ISO/IEC 42001:2023 Artificial intelligence: Management system. https://www.iso.org/standard/42001

17. NIST. Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

18. Osborne Clarke. EU pharma regulators publish annual AI observatory report, 2026. https://www.osborneclarke.com/insights/eu-pharma-regulators-publish-annual-ai-observatory-report

Disclaimer: This white paper is for information only and does not constitute legal, regulatory or technical advice. It reflects public sources available on 5 November 2026. Several sources are drafts or voluntary principles that may change. The framework, risk tiers and roadmap are HERMS planning models, not regulatory requirements. © 2026 HERMS Consultancy DMCC.

Subject index

Select a topic to go to its section in this white paper.

AI risk tiers

Evidence traceability and audit

Generative AI

Governed AI in regulatory affairs

Human decision rights

Human-accountable intelligence

Regulatory AI landscape

Related HERMS resources

Life sciences software for GCC teams

HERMS life sciences solutions

GCC pharmacovigilance obligations for marketing authorisation holders

UAE pharmaceutical product registration

Contact HERMS

Linked references and further reading

https://ai-act-service-desk.ec.europa.eu/en/ai-act/timeline/timeline-implementation-eu-ai-act

https://health.ec.europa.eu/document/download/5f38a92d-bb8e-4264-8898-ea076e926db6_en?filename=mp_vol4_chap4_annex22_consultation_guideline_en.pdf

https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

https://sdaia.gov.sa/en/SDAIA/about/Documents/ai-principles.pdf

https://sdaia.gov.sa/en/SDAIA/about/Files/GenAIGuidelinesForGovernmentENCompressed.pdf

https://www.dlapiper.com/en-at/insights/publications/2025/01/fda-releases-draft-guidance-on-use-of-ai

https://www.doh.gov.ae/-/media/B84FF5DD027846699C51F36C439A580F.ashx

https://www.doh.gov.ae/en/news/the-department-of-health-becomes-region

https://www.ema.europa.eu/en/documents/presentation/presentation-delivering-stakeholders-requirements-ai-l-pinheiro-ema_en.pdf

https://www.ema.europa.eu/en/events/good-manufacturing-practice-multistakeholder-workshop-expert-contributions-artificial-intelligence-guidance-development-annex-22

https://www.ema.europa.eu/en/news/ema-fda-set-common-principles-ai-medicine-development-0

https://www.ema.europa.eu/en/news/harnessing-ai-medicines-regulation-use-large-language-models-llms

https://www.ema.europa.eu/en/scientific-guidelines/use-artificial-intelligence-ai-medicinal-product-lifecycle

https://www.fda.gov/about-fda/center-drug-evaluation-and-research-cder/artificial-intelligence-drug-development

https://www.fda.gov/media/189581/download

https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-use-artificial-intelligence-support-regulatory-decision-making-drug-and-biological

https://www.hermsglobal.com/contact-us

https://www.hermsglobal.com/life-sciences-software-solutions-gcc

https://www.hermsglobal.com/pharmaceutical-product-registration-uae

https://www.hermsglobal.com/pharmacovigilance-obligations-gcc-mah

https://www.hermsglobal.com/solutions

https://www.iso.org/standard/42001

https://www.osborneclarke.com/insights/eu-pharma-regulators-publish-annual-ai-observatory-report