

PDA/FDA Joint Regulatory Conference 2026



14-16 September | Washington, DC

pda.org/pdafda2026
#PDAFDA



Demystifying How Regulators Use AI

Parenteral Drug Association (PDA) | September 2026

Hesha Duggirala, PhD
U.S. Food and Drug Administration

Why This Talk?

- Industry has real questions and some misconceptions about how FDA uses AI
- Goal: transparency, not mystery
- Understanding each other's AI practices builds trust and reduces friction
- Key distinction: FDA's use of AI to assist investigators is very different from industry's use of AI in manufacturing
- This talk answers the questions we hear most often



What Industry Is Asking

Is AI targeting my facility for inspection?

When FDA collects my data, is it safe?

Is AI reviewing my submissions — without a human?

Will FDA scrutinize how I use AI during an inspection?

Does FDA even have guardrails for its own AI use?

How is FDA training its staff to use these tools?



FDA's AI Landscape: What Tools Exist?

Elsa — FDA's internal generative AI tool (agency-wide)

ChatGPT @ HHS and Claude @ HHS — HHS-managed tools available to FDA staff

Each tool has a defined scope, access controls, and approved use cases

MYTH BUST: Document collection is not the same as AI review

- Investigators collect exhibits using a variety of approved methods: paper documents, scanned copies, USB drives, CDs/DVDs, Box, or other approved platforms
- The collection method does not determine whether AI is involved; AI enters the picture only at the review and processing stage

What Is Elsa?

FDA's internal generative AI tool, based on Anthropic's Claude Sonnet and Google's Gemini models

Hosted on AWS GovCloud + Google Cloud Platform (FedRAMP / FISMA High)

Accessible only within the FDA network (no direct internet or intranet connection)

Not trained on submission data or any non-public FDA data

Accepts: text, images, Word, PDF, PowerPoint, CSV files

MYTH BUST: Elsa is not autonomously reviewing your submissions or inspecting your data

What Elsa Can — and Cannot — Do

-  ✓ Summarize long or complex documents
-  ✓ Draft text, emails, and communications
-  ✓ Answer questions based on uploaded files
-  ✓ Assist with image analysis and basic coding tasks
- ✓ ✗ Not connected to the internet or FDA intranet
-  ✗ Does not retain information between chats
-  ✗ Does not make regulatory decisions
- ” ✗ Is not an authoritative source, all outputs must be verified

Governance: The Backbone of Responsible AI Use

- The AI Internal Council (AIIC) is FDA's principal oversight body for **internal AI use**, governing how FDA staff use AI tools to do their work, distinct from FDA's oversight of AI in medical devices or AI used by regulated industry
- Staffed entirely by FDA full-time employees, drawn from all centers, offices, and programs, selected by FDA leadership and center directors, making decisions for FDA, from within FDA
- Validates high-impact use cases before deployment
- Publishes best practices guides and monitors ongoing use
- Aligned with federal AI policy: Executive Order 14179 and OMB M-25-21
- **KEY MESSAGE: Governance is not an afterthought, it is foundational to how FDA uses AI internally**

Upskilling FDA Staff on AI

- All FDA staff receive access to training materials, webinars, and virtual office hours
- Network of 'Elsa Champions' across every center and office supports day-to-day questions
- Core principle: Elsa is assistive, not authoritative, outputs are always verified by domain experts
- Trainings cover: basic Q&A → sophisticated document analysis
- The human expert is always responsible for the decision, not AI
- Bias mitigation: outputs are cross-checked against authoritative sources; generative AI does not make regulatory decisions



How Does AI Use Impact Inspections?

Inspection:

- Investigator prepares using available compliance history, prior observations, and background documents
- On-site: document review/evaluation, interviews, observational walkthroughs, etc.; investigating and evidence gathering for observed deficiencies; drafting 483 observations / closeout.
- Sterile drug inspections often involve large volumes of complex records

AI interface:

- Document collection is logistics, AI enters only at the review stage
- AI tools can assist investigators in organizing regulatory notes, flagging relevant records, or supporting data entry, but **the investigator drives every decision**
- **The human is always in the loop: AI surfaces information; qualified FDA staff evaluate it, identify potential observations, and follow-up.**

Case Study 1: Elsa in the Day-to-Day Inspection Workflow

- Scenario: CSO receives large amount of firm background, history, and administrative information (e.g. Site Master File, etc.)
- Elsa is used to summarize content, identify key data points, and organize information into structured notes
- Elsa may assist with data entry support and organizing observations into a structured format. The investigator authors, reviews, and owns all final work product.
- No AI output is accepted without human review
- **Takeaway: Elsa reduces time spent on administrative burden, freeing investigators to focus on substantive analysis and judgment**

Case Study 2: More Sophisticated Analysis

- Scenario: Investigator needs to identify patterns across a large volume of exhibit data during a complex inspection
- AI-assisted tools can help surface anomalies or inconsistencies at scale. For example, grouping complaint records to identify trends or cross-referencing datasets across exhibits
- Investigator investigates and identifies observations (AI flags, human decides)
- All AI outputs are treated as preliminary and subject to expert verification
- What this is NOT: AI is not issuing findings, writing citations, or making compliance determinations
- Human expertise and judgment are the foundation of every inspection outcome

What AI Does NOT Do in FDA Inspections

 **X** Does not autonomously select or target facilities for inspection

 **X** Does not make or issue regulatory decisions

 **X** Does not replace investigator judgment or expertise at any stage

 **X** Is not the decision-maker in any regulatory action

 The regulatory decision is always made by qualified humans, full stop

AI in Your Facilities




FDA and global regulators are watching how industry uses AI in CGMP settings — and scrutiny is increasing



FDA and MHRA: the concern is not whether AI is used — it is whether outputs are **accurate, verifiable, technically reviewed, and prepared under appropriate human oversight**



FDA does not prohibit AI use — but human oversight, review, and approval of AI-generated outputs is required on both sides of the inspection equation



April 2026 FDA warning letter: A manufacturer was cited for improper reliance on AI without adequate human oversight — AI-generated SOPs, batch records, and specifications lacking human review constitute potential CGMP violations



Warning signs regulators look for: generic language, non-existent citations, lack of facility-specific detail, inconsistent terminology, and verbose responses that don't address the actual deficiency



June 2026 MHRA guidance blog: MHRA has identified AI-drafted inspection responses containing fabricated regulatory references, citations of non-existent guidance, and responses to serious deficiencies that failed to address the underlying problems — in one case, review time increased from ~4 hours to over 20 hours due to an inaccurate, AI-generated response to a patient safety deficiency



The principle is clear: AI writing a response or SOP is not the failure — an inaccurate response or SOP is the failure. AI may be the root cause, but accountability remains with the organization

What FDA Expects When It Sees AI in Your Quality Systems

- Clear documented purpose for each AI tool, including its **context of use**: what it does, what data it acts on, and what decisions it informs
- Validated performance appropriate to that context of use with documentation integrated into quality system records and SOPs
- Controls on training data, input data, and model outputs with traceability throughout
- A lifecycle plan: managing model drift, updates, and changes under change control
- Human review and approval of AI-generated outputs, always documented
- AI tools used in GxP activities must be validated for their intended use and fit for the specific regulatory context in which they operate

Emerging Regulatory Landscape

- CDER 2026 Guidance Agenda: 'AI/ML Quality Considerations in Pharmaceutical Manufacturing' — forthcoming
- FDA Jan 2025 draft guidance: Seven-step credibility assessment framework for AI in regulatory submissions. Review and comment period on this draft has closed.
- FDA / EMA joint guiding principles for AI in drug development — transatlantic harmonization underway*

The regulatory environment is evolving rapidly — engagement with public dockets is strongly encouraged


* FDA and EMA. Guiding Principles of Good AI Practice in Drug Development. January 2026.

Data Integrity in the Age of AI

AI does not eliminate data integrity obligations



Existing CGMP rules apply. Cloud, IoT, and AI-generated records fall under FDA's 2018 Data Integrity Q&A guidance. ALCOA+ principles apply whether a human or an algorithm created the record.



Investigators are asking. Expect scrutiny of provenance, model oversight, and who reviewed AI-generated outputs. PDA Week 2026 flagged AI-assisted deviation management and data integrity hotspots as active inspection themes.



Your best defense: proactive documentation. Record what your AI tools do, how they are validated, and how outputs are reviewed and approved before an investigator asks.

How Can Industry Help FDA?

- Engage with FDA's public AI communications, Q&A documents, and guidance dockets
- Ask questions, FDA wants dialogue, not confusion or speculation
- Be transparent during inspections about what AI tools you use and how they are governed
- Think of your own AI governance structure: validation, oversight, documentation
- Shared goal: AI that supports public health without compromising regulatory integrity



What 'Helping FDA' Looks Like in Practice

- Don't assume AI use is invisible to investigators, it isn't
- Don't use AI as a compliance shortcut, it will attract scrutiny
- Don't rely on a single reviewer to validate all AI outputs, the expertise of the reviewer must match the nature of the output: a QC professional reviewing AI-assisted SOP language, a statistician reviewing AI-assisted data analysis, a manufacturing expert reviewing AI-generated batch record summaries
- Do document your AI use cases, validation approach, and human oversight steps, including **who reviewed what, and what qualifications they brought to that review**
- Do engage with FDA's dockets, public meetings, and comment periods on AI guidance
- Do bring your questions to forums like this one, open dialogue makes better policy

Key Takeaways

1. FDA uses AI to assist investigators, not to make decisions
2. Governance and human oversight are non-negotiable at FDA
3. Industry's use of AI in CGMP settings is under increasing regulatory scrutiny
4. Transparency and documentation are your best tools (on both sides)
5. The human-in-the-loop principle is universal: it applies to FDA and to you
6. This is a conversation FDA wants to have with industry (keep talking)