FDA Elsa agency-wide generative AI tool

IMPLEMENTED 2025–2026 · GOVERNMENT AND PUBLIC SERVICES · US · BAND D · CLAIMED ONLY

VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

The US Food and Drug Administration launched Elsa, a generative AI assistant for all staff, on 2 June 2025, and has since added agentic features and a consolidated data platform. The agency reports voluntary use by more than 70 per cent of staff and says Elsa speeds reviews; staff quoted by CNN and NOTUS reported invented studies and guidance, which the agency disputed.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Generative AI assistant with agentic features: Elsa.

Implemented: 2025–2026. Grade: band D, Task support, Claimed only.

Headline result, as reported: Agency reports voluntary use by more than 70 per cent of staff; staff reported fabricated citations, which the agency disputed

1 · The case

The facts, as reported by the 5 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

US Food and Drug Administration. US federal regulator within the Department of Health and Human Services, responsible for human and veterinary drugs, vaccines, medical devices, food, cosmetics and tobacco. Elsa is available to all staff, from scientific reviewers to investigators. Canvas block: 5. People and skills. Also 4. Operations and process, 8. Governance and risk.

Business challenge

The challenge. By the agency's account, FDA staff handle large volumes of documents in scientific reviews, inspections and safety assessments. Its commissioner described a pilot with scientific reviewers, then set a timeline to scale AI agency-wide by 30 June 2025.

Who implemented it

Google Cloud (Model or cloud provider): Hosts the FedRAMP High environment in which Elsa 4.0 is built, per the FDA's May 2026 release.
In-house (In-house): Built and continues to develop Elsa with agency reviewers.
The model provider is reported by one reviewer and is not confirmed by the FDA.

Technology (SEAM pillar E)

SEAM types: T1 Chat assistants, T7 AI agents. Deployment: Private cloud. Vendors and products: Google Cloud, Elsa, Elsa 4.0, HALO data platform.

Where it was built in: Used by FDA staff across centres for reading, writing and summarising. From May 2026 integrated with HALO, which consolidates more than 40 application and submission data sources, so that staff can query data without uploading documents in each chat.

What changed

Before: staff read, compared and summarised documents themselves. After, by the agency's account: staff use Elsa to summarise adverse events, compare labels, generate code for databases and accelerate protocol reviews, and from May 2026 Elsa sits on top of the agency's data rather than staff bringing data to it. How much review time changed: Not reported. [1] [2] [3] [4] [5]

What it cost

The FDA says the launch was under budget. A leaked internal slide reported by Gizmodo put the cost at $12,000 in the first four weeks. Total cost: Not reported.

What it saved

Not reported. The FDA says Elsa is shortening the time needed for scientific evaluations and that HALO and Elsa 4.0 are expected to create cost savings, but gives no figures in the releases opened. [1] [2] [3] [4] [5]

Governance and control

The FDA says staff are involved at every stage and that human experts verify inputs, analytic processes and outputs, that agentic use is optional with built-in human oversight, and that models do not train on input or industry-submitted data. HHS told NOTUS the agency had put in guardrails, disclaimers and trainings. No external review of Elsa was found.

S2 · Accountable owner · Claimed
The FDA names its Chief AI Officer, Jeremy Walsh, in its releases, and says its Office of Digital Transformation has rolled out improvements and new use cases. [1] [3]

A1 · Safe use · Claimed
HHS told NOTUS the FDA had put in guardrails, disclaimers and trainings for staff; the FDA says models do not train on data submitted by regulated industry. [1] [5]

A2 · Human in the loop · Claimed
The FDA says staff are involved at every stage and human experts verify inputs, analytic processes and output implementation; staff quoted by CNN said unreliable output needs double-checking. [3] [4]

A3 · Monitoring and assurance · Not reported
No published accuracy evaluation, audit or error monitoring was found; the 70 per cent usage figure is internal agency data. [2]

A4 · Vendor AI · Claimed
The FDA says Elsa is built in a secure cloud environment and does not train on regulated-industry data; a reviewer told NOTUS the agency limits the model provider's access to agency files. [3] [5]

A5 · Agent controls · Claimed
The FDA says agentic use is entirely optional and the tools have built-in guidelines including human oversight. [2]

A6 · Standards and regulation · Not reported
No standard or framework applied to Elsa is named in the sources opened.

What worked, what did not, what is unproven

WORKED

– Staff use it voluntarily: more than 70 per cent, according to internal agency data cited by the FDA in December 2025.

– Three of the six current and former employees who spoke to CNN described it as helpful for tasks such as meeting notes and summaries (per Gizmodo).

– The agency has kept developing it, moving from the June 2025 launch to agentic features in December 2025 and Elsa 4.0 in May 2026.

DID NOT

– Three employees told CNN that Elsa produced nonexistent studies and misrepresented real research, and one said "It hallucinates confidently" (per Gizmodo).

– One FDA reviewer told NOTUS that Elsa appeared to quote FDA guidance that did not exist, and that they stopped using it to draft letters to drug companies.

– Staff testing with basic questions, such as how many drugs in a class were approved for children, produced wrong answers, per the CNN report as relayed by Gizmodo.

– The agency's spokesperson said CNN had mischaracterised and taken the information out of context; HHS said large language models inherently hallucinate.

UNPROVEN

– Whether Elsa shortens reviews: the FDA says so, and the sources opened give no measurement.

– How often Elsa is wrong: the sources are anonymous staff accounts, and no error rate is published.

– Which model powers Elsa: the FDA's releases name only the cloud environment, and no model provider is confirmed.

– Whether the original rollout problems persist in Elsa 4.0: the FDA's spokesperson said critics had not used the current version.

Timeline

Jul 2025 CNN reports staff accounts of invented studies; NOTUS reports similar staff accounts; FDA disputes the CNN report [4]

Dec 2025 FDA deploys optional agentic AI capabilities to all employees and reports Elsa use by more than 70 per cent of staff [2]

May 2026 FDA launches Elsa 4.0 and the HALO data platform [3]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: A. Grade: D, Task support. Proof: Claimed only. Level of change: Task support.

Task support (Level 1, band D). Elsa is an optional assistant with human verification and has been extended with agents and a data platform. The agency claims transformation; the sources opened show no measured change in how reviews are done.

The FDA describes Elsa as an optional assistant for reading, writing and summarising, and says its commissioner described most use as organisational tasks. That is task support. The agency says it is accelerating reviews and has added custom agents and a data platform, and its launch release speaks of transforming operations through AI, so the agency claims a larger change. The sources opened give no measured change in review times or in the way reviews are done, so the transformation is claimed and unproven.

The SEAM view and lessons learnt

A regulator rolled out a generative AI assistant to all staff quickly and has kept extending it. Two accounts sit side by side: the agency's own, with usage and speed claims, and staff accounts of invented citations. Neither is tested in the sources opened, and no accuracy evaluation has been published. For a mid-sized or larger organisation, the transferable point is that voluntary adoption figures say how many people use a tool, and say little about whether its outputs are right. Citation and quotation checks need to be built into the workflow before use spreads.

1. Publish an accuracy test on real tasks before and after each version.
2. Make staff verify every citation and quotation the tool produces.
3. Report adoption and error rates together.

Evidence

Score 1 of 4. Company account: FDA press releases, plus press reports of anonymous staff accounts. No independent evaluation, audit or published error rate. Usage figures are internal agency data.

Evaluation method. No evaluation reported. FDA releases give usage and feature claims; CNN (as relayed by Gizmodo) and NOTUS report accounts from current and former staff.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/fda-elsa-generative-ai-tool

[1] US Food and Drug Administration: FDA Launches Agency-Wide AI Tool to Optimize Performance for the American People (2 June 2025)
https://www.fda.gov/news-events/press-announcements/fda-launches-agency-wide-ai-tool-optimize-performance-american-people

[2] US Food and Drug Administration: FDA Expands Artificial Intelligence Capabilities with Agentic AI Deployment (1 December 2025)
https://www.fda.gov/news-events/press-announcements/fda-expands-artificial-intelligence-capabilities-agentic-ai-deployment

[3] US Food and Drug Administration: FDA Expands AI Capabilities and Completes Data Platform Consolidation (6 May 2026)
https://www.fda.gov/news-events/press-announcements/fda-expands-ai-capabilities-and-completes-data-platform-consolidation

[4] Gizmodo: FDA's New Drug Approval AI Is Generating Fake Studies: Report (23 July 2025, relaying CNN)
https://gizmodo.com/fdas-new-drug-approval-ai-is-generating-fake-studies-report-2000633153

[5] NOTUS: 'They Have To Figure Out If They Can Trust The Software': Inside the FDA's AI Push (29 July 2025)
https://www.notus.org/health-science/fda-ai-push-kennedy-elsa-makary

US Food and Drug Administration website: https://www.fda.gov

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 26 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/fda-elsa-generative-ai-tool

VERSION HISTORY

v1.0 15 November 2026 First published