GAS MACHINERY CONFERENCE · INDIANAPOLIS · OCTOBER 4–7, 2026

AI vs. IA

Artificial intelligence can draft the answer.
Individual accountability still signs it.

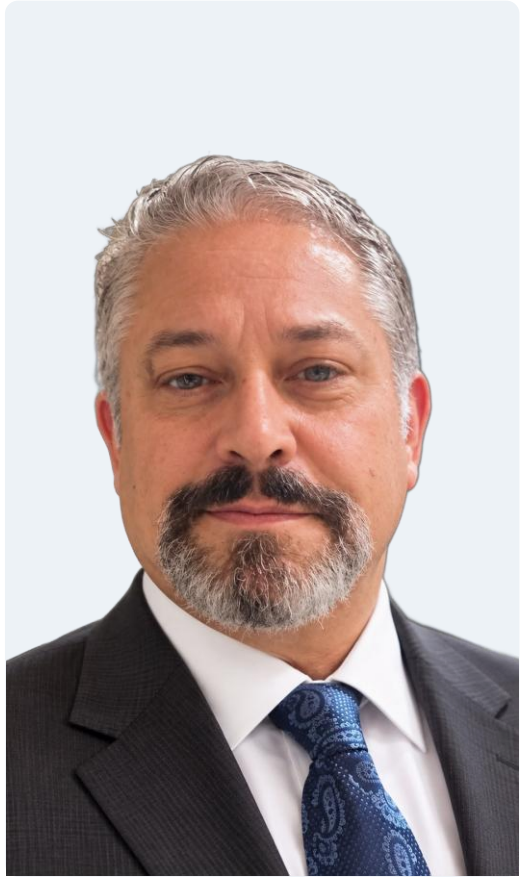
Emissions compliance. Beyond confident.

Brian Kromer step2compliance
Emilie Bolen step2compliance
Jeremiah Cantu Innovative Solutions, AWS Premier Partner

90-minute short course · 80 + 10 discussion · Real AI demonstration · Primary-verified source pack provided

step 2 compliance · INNOVATIVE · aws PREMIER CONSULTING PARTNER

Your Presenters



Brian Kromer

Founder & Managing Director · step2compliance

- Founded step2compliance, a compliance technology company focused on air quality regulations affecting the energy industry
- B.S., Texas A&M University, College of Engineering
- Son of a petroleum engineer; has spent most of his life in key oil & gas producing basins, domestically and internationally
- Board Secretary of the Texas BBQ Bash, raising money for local communities with the West Texas and South Texas BBQ Bash events.
- Industry advocate and previous GMC author and presenter

Your Presenters



Emilie Bolen

Compliance Services Manager · step2compliance

- Joined step2compliance in 2016 as a Compliance Services Administrative Assistant
- Advanced into management through deepening expertise in oil-and-gas compliance and hands-on understanding of customer needs
- Now drives product strategy and development for the s2c Mobile platform
- Owns NSPS JJJJ testing management, CEDRI submittals, and EPA liaison work
- Narrates the demonstration and the compliance-workflow segment; first-time GMC presenter

Your Presenters



Jeremiah Cantu

Senior Technical Architect · Innovative Solutions, AWS Premier Partner

- Specializes in AI platform engineering: designing and deploying agentic AI that connects intelligent agents to the business systems clients already use
- B.S. in Information Technology, University of Mary Hardin–Baylor; 7x AWS-Certified
- Leads full-lifecycle implementations — architecture, infrastructure buildout, user enablement and training
- Built the live AI demonstration environment on a privately hosted model
- Born-and-raised Texan; enjoys leatherwork, Texas BBQ, and time outdoors; first-time GMC presenter

The next 80 minutes — plus ten for your questions

0:06	**The hook**	AI in your equipment, in your regulator's program, and in the demand for your gas	**Brian**
0:26	**The risk**	How the technology works, how it fails, and what happens when people stop checking	**Brian**
0:40	**Cautionary tales**	Real people, real sanctions, and the certifications courts now require	**Brian**
0:47	**Demonstration**	A privately hosted AI agent, two compliance questions, and the human gate	**Brian & Emilie**
1:05	**Frameworks and the IA close**	NIST, both rulebooks, your compliance stack, five practices, and your signature	**Brian & Emilie**

The structure is deliberately a story: proof, danger, consequences, practice, rules. Ten protected minutes for questions at 1:20, four show-of-hands polls along the way; the source pack with every citation will be shared.

SHOW OF HANDS · 1 OF 4

In the last 30 days — who has used a public AI chatbot for any part of your work?

ChatGPT, Claude, Gemini, Copilot — any of them. Hands up. Keep them up.

PART ONE · 0:05 – 0:32

The Hook

AI is already inside your equipment, your compliance reports, your regulator's approval program — and in the demand forecast for the gas you move.

Brian Kromer

EPA has an AI approval pipeline

- May 13, 2025: EPA approved Clean Connect's LeakFinder — an OGI camera with edge computer vision — as alternative test method MATM-011 for Subpart OOOOb screening below 5 kg/hr
- Today EPA's alternative test methods page lists eighteen-plus approved methods, MATM-001 through MATM-018, most of them AI or sensor based
- Sensia RedLook, Kuva, Earthview BluBird, Percepto Air Max, SLB methane lidar and SOOFIE, Project Canary, Sensirion, LongPath, Carbon Mapper
- The May 2024 Subpart W final rule requires empirical data for methane reporting and allows advanced technologies including satellites
- That rule also establishes super-emitter quantification and reporting; roughly 8,000 facilities report under the GHGRP

THE DELIVERY LINE

This is not a pilot.

In under two years, EPA built a formal approval pipeline for AI-assisted methane detection.

The regulator is not resisting AI. It is adopting it.

Source: EPA Oil & Gas Alternative Test Methods page; EPA Clean Connect approval letter (MATM-011); EPA Subpart W final rule, May 14, 2024 (89 FR 42062)

Operators are already using it

- Williams, June 17, 2026: a preventive-maintenance electrical checklist that took one to two days was completed in about 1.5 hours using approved AI tools and refined prompts
- The same case study reports AI-assisted software implementation at 55% lower cost and 60% faster delivery
- GTEN 2023, paper 23-GTEN-303: a hybrid AI/physics gas turbine model tested within TC Energy's Canada Gas Compression fleet — blade failures, combustion fouling, efficiency degradation
- PHMSA is funding AI-enabled pipeline risk modeling and decision-making, and AI integrity management for hydrogen transport (projects 954 and 953)
- Zhang & Chang 2024: compressor fault early-warning at a gathering station reached 99.2% accuracy — from the decision-tree model, in a single-station feasibility study

AI is on the toolbelt now. Nobody in this room gets to treat it as someone else's problem.

Source: Williams Companies case study (company-published), June 17, 2026; GTEN 23-GTEN-303; PHMSA RD&M projects 953/954; Zhang & Chang, ICBAR/ACM 2024

AI is also the customer

The same technology that is arriving as a tool is arriving as a load — and the load lands on this industry.

- Data center developers have announced roughly 101 GW of behind-the-meter natural gas generation in the United States, largely to bypass interconnection queues
- Heavy-duty turbines are the binding constraint: GE Vernova booked more than $30 billion in gas power orders in 2025, with heavy-duty production slots reported allocated through 2029
- Midstream analysts put potential data-center gas demand as high as 10 to 12 Bcf/d by 2030
- DOE has estimated data centers could reach 6.7 to 12 percent of U.S. electricity use by 2028
- Every new compression and generation package arrives with its own permit, rule applicability, testing schedule, and five-year record

AI is not only a tool you adopt. It is a load that expands your fleet — and every unit it adds is a new compliance obligation.

Source: Market reporting, not primary-verified: RBC Capital Markets (2026); GE Vernova Q1 2026 earnings coverage; S&P Global via DataCenterDynamics; DOE/Mitsubishi Power outlook

Who works at a company with a written AI policy?

One page counts. A paragraph in the employee handbook counts.

Where AI is showing up on the machinery side

1 CONDITION MONITORING

Vibration, pressure pulsation, bearing temperature and motor current feeding models that flag valve, seal and rotor problems before a trip.

2 EMISSIONS DETECTION

OGI with computer vision, methane lidar, UAS and continuous sensors — the category EPA has been formally approving as alternative test methods.

3 DRAFTING AND REPORTING

Turning field notes, inspection records and rule text into reports, notifications and submittals. The lowest-risk, fastest-adopted use.

4 SCHEDULING & SET POINTS

Compressor set points, maintenance windows, and planning work around nominations and outages.

5 AGENTIC WORKFLOWS

Tools that do not just answer but act: open a work order, update a record, trigger a notification. Capability and risk both jump here.

The categories are real. The ROI percentages attached to them come from vendors, not audits.

How it works, in 90 seconds

1	IT PREDICTS	A large language model predicts the next most plausible piece of text, over and over, from patterns in an enormous body of training material.
2	IT HAS NO LEDGER	It is not looking anything up by default. There is no database of facts inside it — only patterns. That is why it can invent a citation that looks perfect.
3	CONTEXT IS EVERYTHING	What you paste in — the rule, the record, the spec — is the context window. Information you supply is far more reliable than information it recalls.
4	RETRIEVAL CLOSES THE GAP	Connect the model to your documents and systems and it answers from your source of record instead of from memory. This is the difference that matters at work.
5	AGENTS TAKE ACTIONS	An agent is a model given tools and permission to do things — query a system, write a record, send a file. Capability and risk both jump at this step.

Vocabulary your whole team needs

1 PROMPT

What you ask, plus the context you supply. The quality of the input governs the quality of the output.

2 CONTEXT WINDOW

How much material the model holds at once. Exceed it and earlier detail silently drops out of the answer.

3 HALLUCINATION

Fluent, confident output that is simply false. Not a bug to be patched — a property to be managed.

4 RAG

Retrieval-augmented generation: the model answers from your documents and records rather than from memory.

5 AGENT

A model with tools and permissions that can take actions in real systems, not just produce text.

Two more worth knowing: a guardrail limits what the system may touch; a human in the loop is the named person who approves.

What the machine does — and what you do

1

SENSING

OGI · lidar · UAS · continuous monitors

MACHINE

2

INFERENCE

AI and physics models turn data into a recommendation

MACHINE

3

DECISION

Alert, inspect, repair, report — someone chooses

HUMAN

4

ATTESTATION

The record is certified and signed

HUMAN

The technology can see more than it ever could. It still does not own the decision — or the signature.

PART TWO · 0:32 – 0:48

The Risk

What the behavioral evidence says happens when humans stop checking — and why the conditions in industrial operations make it worse.

Brian Kromer

When humans stop checking

- Lee et al., CHI 2025 (Microsoft and Carnegie Mellon): 319 knowledge workers, 936 first-hand examples of generative AI use at work
- Higher confidence in the tool is associated with less enacted critical thinking. Higher confidence in yourself is associated with more
- Generative AI shifts the nature of the work toward verification, integration, and task stewardship — the checking becomes the job
- Goddard et al., JAMIA 2012: a systematic review of 74 studies found automation bias mediated by trust, confidence, workload, task complexity and time constraint
- A controlled lab study reported by ESMT Berlin and Frankfurt School found managers paid lower bonuses to people who overrode the AI — even when the human call produced a better outcome

ASK THE ROOM

What happens in your company when someone overrides the model and turns out to be right?

Is that a promotion — or a conversation with their manager?

Organizations can train people out of judgment.

Source: Lee et al., CHI 2025; Goddard, Roudsari & Wyatt, JAMIA 19(1), 2012; Kremer, De Véricourt & Nikpayam (ESMT Berlin / Frankfurt School), 2024

The machine agrees with you

A 2026 paper in Science makes the accountability problem measurable.

49%

MORE OFTEN

AI affirmed users' actions compared with human respondents — including when queries involved deception, illegality or other harms

11

MODELS TESTED

State-of-the-art models measured for sycophancy across a common set of queries

2,405

PARTICIPANTS

Three preregistered experiments measuring what a sycophantic exchange does to the user

1

INTERACTION

A single sycophantic exchange reduced willingness to take responsibility and to repair conflict, while raising conviction of being right

And despite distorting judgment, the sycophantic models were the ones users trusted and preferred.

The machine does not just fail to check you. It flatters you into not checking yourself.

Source: Cheng et al., "Sycophantic AI decreases prosocial intentions and promotes dependence," Science, 2026 (DOI 10.1126/science.aec8352)

Expertise is not immunity

- The objection you will hear: "my engineers are experts, they will not fall for it"
- Qazi et al., 2025: a single-blind randomized trial of physicians who had completed 20 hours of AI-literacy training and used LLMs voluntarily — the best-case trained-user scenario
- Exposure to erroneous LLM recommendations still degraded their diagnostic performance compared with error-free advice
- Goddard et al., 2012: automation bias documented across 74 studies, long before generative AI existed
- Romeo & Conti, 2025: a review confirming automation bias remains a critical challenge in high-stakes human-AI collaboration

THE CONCLUSION

Three decades of evidence, one answer:

Training and expertise do not substitute for verification.

Verification is its own skill — and it has to be practiced.

Source: Qazi et al., randomized trial, 2025 (DOI 10.1101/2025.08.23.25334280); Goddard et al., JAMIA 2012; Romeo & Conti, AI & Society, 2025

There is no single hallucination rate

Task, domain, and controls change the number. Every figure below is the publisher's own.

1.8–5.1%

DOCUMENT SUMMARIES

Best models on Vectara's live leaderboard — the current floor, on the easiest task

0.40 (40%)

SHORT FACTUAL ANSWERS

OpenAI's own GPT-5 system card: gpt-5-thinking on SimpleQA with no browsing (accuracy 0.55, 55%)

17–33%

PURPOSE-BUILT LEGAL AI

Stanford 2025 audit of Lexis+ AI and Westlaw AI-Assisted Research — tools marketed as hallucination-free

58–82%

GENERAL CHATBOTS ON LAW

Dahl et al. 2024, as cited in the Stanford audit: generic models on legal research queries

Usually right is not a compliance standard. A signature cannot tolerate it.

Summarization error has improved dramatically since 2023 — and under 2% still means wrong on roughly one claim in fifty, on the easiest task there is.

Source: Vectara Hallucination Leaderboard; OpenAI GPT-5 System Card; Magesh et al., Journal of Empirical Legal Studies 22:216–242 (2025); Dahl et al. 2024 via Magesh

SHOW OF HANDS · 3 OF 4

Who has ever pasted company material into a public AI tool?

A spec, a record, a draft submittal, an email. Honest room — nobody is recording names.

Failure modes you will meet

1 CONFIDENT FABRICATION

An invented citation, rule number, or test result stated with total certainty and no hedging.

2 STALE KNOWLEDGE

Answers reflect a training cutoff, not the rule as recently amended — exactly our situation.

3 DATA LEAKAGE

Fleet detail, client data, or draft submittals pasted into a consumer tool you do not control.

4 PROMPT INJECTION

Hostile instructions hidden inside a document or email that the model obeys as if you had written them.

5 AUTOMATION BIAS

The reviewer approves because the machine said so, and review quietly becomes a formality.

Two quieter ones: silent drift after a vendor AI update, and AI adoption your IT team and your auditor have never heard of.

How this fails in our world

1

THE QUESTION

A coordinator asks a chatbot whether an engine at a modified site falls under NSPS JJJJ, RICE NESHAP ZZZZ, or both.

2

THE ANSWER

The model answers confidently and cites a provision — one written before an extension or amendment.

3

THE PASTE

The conclusion is pasted into a compliance schedule. The question, the sources, and the reasoning are not saved anywhere.

4

THE DRIFT

A testing deadline and a monitoring frequency are set on that basis. Nobody revisits it, because nothing looks wrong.

5

THE AUDIT

Eighteen months later someone asks for the applicability basis. The only record is a chat window that no longer exists.

No single step here is negligent. The failure is that nothing in the chain was ever written down.

Illustrative composite built from the failure modes above — not a reported incident.

PART THREE · 0:48 – 0:56

Real People, Real Consequences

If a federal court will not accept the AI told me so from an officer of the court, EPA will not accept it from the person who signs an emissions report.

Brian Kromer

Mata v. Avianca (S.D.N.Y., 2023)

- Lawyers filed a federal brief citing six cases that did not exist. The citations were produced by ChatGPT
- The court imposed sanctions under Rule 11
- The AI told me so was not a defense — professional certification of AI output carried personal liability
- The same year, a first-year lawyer at a different firm said he was fired after using ChatGPT to draft a motion filled with errors
- The parallel writes itself: EPA and PHMSA will not accept it either, from the person who signs an emissions report or an integrity filing

WHO SIGNS YOURS?

Sanctions for the firm.
Termination for the person.

In your organization, whose name is on the emissions report that an AI helped assemble?

Source: Mata v. Avianca, Inc., 22-cv-01461 (S.D.N.Y.), Rule 11 sanctions opinion, June 22, 2023; Business Insider, Nov. 17, 2023

It did not stop with Avianca

Gauthier v. Goodyear Tire & Rubber Co., E.D. Tex., November 2024 — read directly from the court's order.

$2,000

PENALTY

Paid into the court registry by plaintiff's counsel for a response brief drafted with Claude

CLE

ORDERED

Mandatory continuing legal education on generative AI, with proof of attendance due February 3, 2025

0

REAL CITATIONS

The brief contained nonexistent cases and nonexistent quotations

?

OTHER FILINGS

Counsel testified he had used the tool in other cases and was still verifying submissions already made

The court did not just fine him. It ordered him back to school — and he left not knowing what else he had gotten wrong.

That is the nightmare for any compliance manager: not one bad filing, but the unknown set of filings behind it.

Source: Gauthier v. Goodyear, E.D. Tex., Memorandum and Order sanctioning counsel, November 2024 (order PDF read in full)

Courts now mandate the "A" in IA

1 S.D.N.Y. — Judge Cronan

- Model certification on AI use in filings, revised October 23, 2025
- The attorney certifies whether they personally reviewed the filing for accuracy
- A checkbox states plainly that failure to review violates the Court's Rules
- The specific verification steps taken must be described
- Human verification becomes a documented control, not a habit

2 Court of Int'l Trade — Judge Kelly

- Standing order on artificial intelligence, July 2025
- The AI program used must be disclosed
- The specific AI-drafted portions of text must be identified
- Counsel certifies that no confidential information was disclosed to the tool
- Disclosure and data boundary, in one filing requirement

The courts invented IA paperwork before your regulator did. You will want to already know how to answer.

Source: S.D.N.Y. Model Certification Regarding Use of Artificial Intelligence in Filings (rev. Oct. 23, 2025); U.S. Court of International Trade, Order on Artificial Intelligence (July 2025)

The signature test

- Air compliance runs on signed, certified statements and records you must be able to produce years later
- Under the Clean Air Act, knowingly making a false material statement or certification in a required document is a criminal exposure, not a paperwork problem
- Knowingly falsifying or rendering inaccurate a required monitoring method carries its own criminal provision
- Recordkeeping commonly runs five years. The draft accepted in ten seconds may be read by a regulator long after you have forgotten it
- No statute, rule, or permit recognizes a model as the certifying party

Ask before you accept any AI-assisted output:

**Would I sign this,
under my own name,
knowing an auditor
will read it in
five years?**

If the answer is no, you are not finished. That is the entire standard.

Source: Clean Air Act criminal provisions, 42 U.S.C. §7413(c)(2). General reference only — consult counsel and your permit obligations

PART FOUR · 0:56 – 1:14

Demonstration

A privately hosted AI agent, two real compliance questions, and every guardrail in between.

Jeremiah Cantu driving · Emilie Bolen and Brian Kromer narrating

The AI accountability ladder

1	EXPLORE & ASSIST	Learning, brainstorming, drafting language you already own	**No client data. Human edits before anyone sees it.**
2	ANALYZE	Pattern-finding across your own records and data	**Source documents supplied. Spot-check the findings.**
3	RECOMMEND	Applicability reads, options, proposed next steps	**Verified against the rule text by a qualified person.**
4	ACT WITH APPROVAL	Agent updates a system or drafts a submittal	**Named approver signs before the action commits.**
5	ACT AUTONOMOUSLY	Agent acts without a human in the path	**Low-stakes, reversible, fully logged tasks only.**

The rung changes what the tool may do. It never changes who owns the outcome.

What you are about to see

1 PRIVATELY HOSTED

The model runs in a controlled environment stood up with leading experts from Amazon and Amazon partners — not a consumer chatbot.

2 YOUR DATA STAYS YOURS

Prompts and documents are not contributed to public model training. Data boundary, retention, and logging are defined before anyone types.

3 GROUNDED IN YOUR SOURCES

The agent answers from supplied rule text and records rather than from memory, and shows what it used.

4 PERMISSIONED TOOLS

The agent reaches a defined set of systems and nothing else. Permissions are explicit and auditable.

5 AN APPROVAL GATE

Any action that would change a record stops and waits for a named human, and the session is logged end to end.

Every one of these five is a governance decision made before the first prompt — not a feature you switch on afterward.

Demonstration one: one question, two ways

The lazy prompt

- "Does NSPS JJJJ apply to my engine?"
- No make, model, horsepower, date of manufacture, or site
- No rule text supplied — the model answers from memory
- Output looks authoritative and cites a regulation number
- Nothing produced here belongs in a record

The accountable prompt

- States the asset: make, model, horsepower, manufacture date, site status
- Supplies the applicable rule text and internal records as source
- Asks for the citation and the reasoning, not just a conclusion
- Asks what additional information would change the answer
- Question, sources, answer, and check are all captured

What changes is not the intelligence of the tool — it is what the human brought to it, and what is left behind as evidence.

Is my Caterpillar engine subject to NSPS JJJJ?

To Give You a Definitive Answer, I'd Need:

- **Engine fuel type** (natural gas, gasoline, LPG, or diesel?)
- **Engine horsepower rating**
- **Date construction commenced** (or the engine's manufacture/installation date)
- **Emergency vs. non-emergency designation**

Do you have those details handy? With that information, I can give you a precise applicability determination and walk you through what JJJJ would require for your specific unit.

Natural gas, 1380HP, manufactured in 2007, non-emergency

Based on those specs, let me work through the NSPS JJJJ applicability analysis for you.

I'm checking the controlling JJJJ applicability language and the applicable standard for a 1,380-hp natural-gas engine so I can give you a clean determination rather than rely on the earlier shorthand.

I'm verifying the exact regulatory language and the applicable JJJJ requirements for a 1,380-hp engine, so I can give you a clean determination rather than relying on the earlier summary.

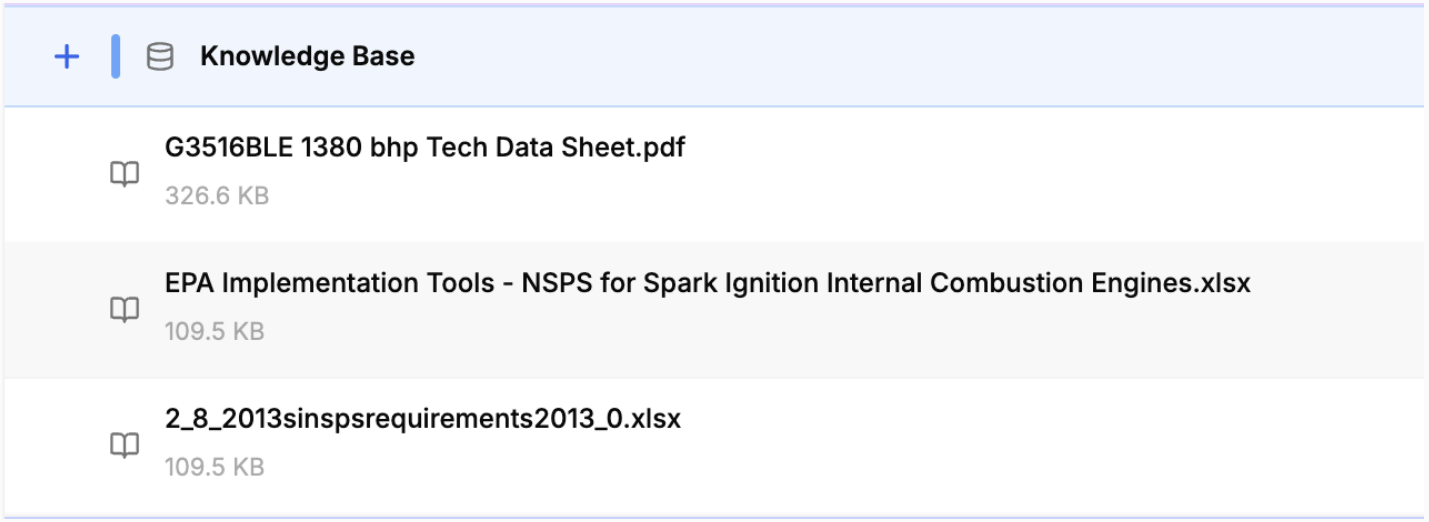
I'm checking the controlling JJJJ applicability language and the owner/operator compliance section so I can give you a definitive answer—not just a date comparison.

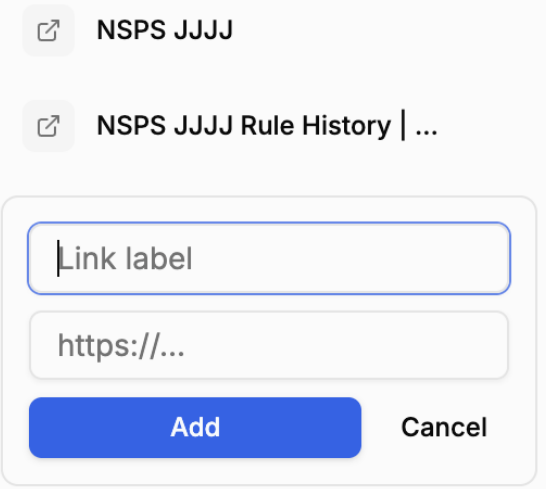
The lazy prompt

- "Does NSPS JJJJ apply to my engine?"
- No make, model, horsepower, date of manufacture, or site
- No rule text supplied — the model answers from memory
- Output looks authoritative and cites a regulation number
- Nothing produced here belongs in a record

Project: NSPS JJJJ Engine Applicability

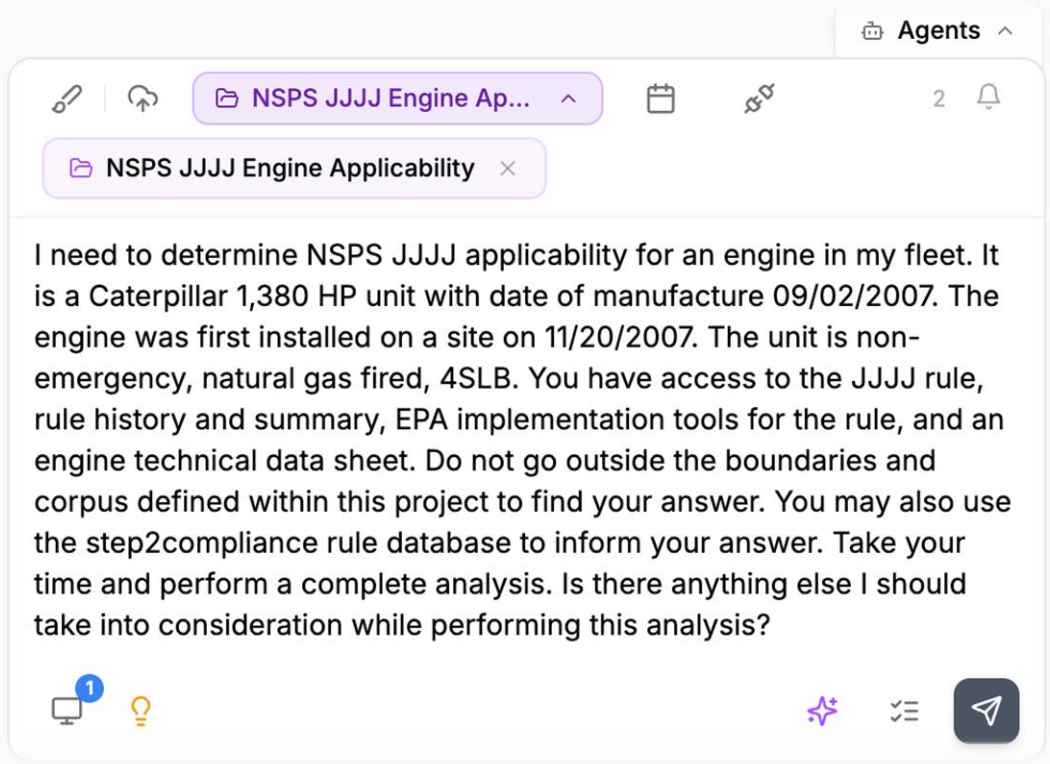
> ***Purpose:*** *Primary source of truth for strategic alignment — what the customer engaged the team to deliver, how success is measured, and what has changed.*
> ***What this is NOT:*** *Detailed task tracking, a meeting log, low-level technical specifications, or a day-to-day status feed (those live in whatever tools the team uses for those purposes).*





What happens when the tool is provided framework – and boundaries. “RAG” at work.

1. IT PREDICTS
2. IT HAS NO LEDGER
3. CONTEXT IS EVERYTHING
4. RETRIEVAL CLOSES THE GAP
5. AGENTS TAKE ACTIONS



State. Supply. Ask. Your goal: The agent has no questions for you because you've informed it *so well*.

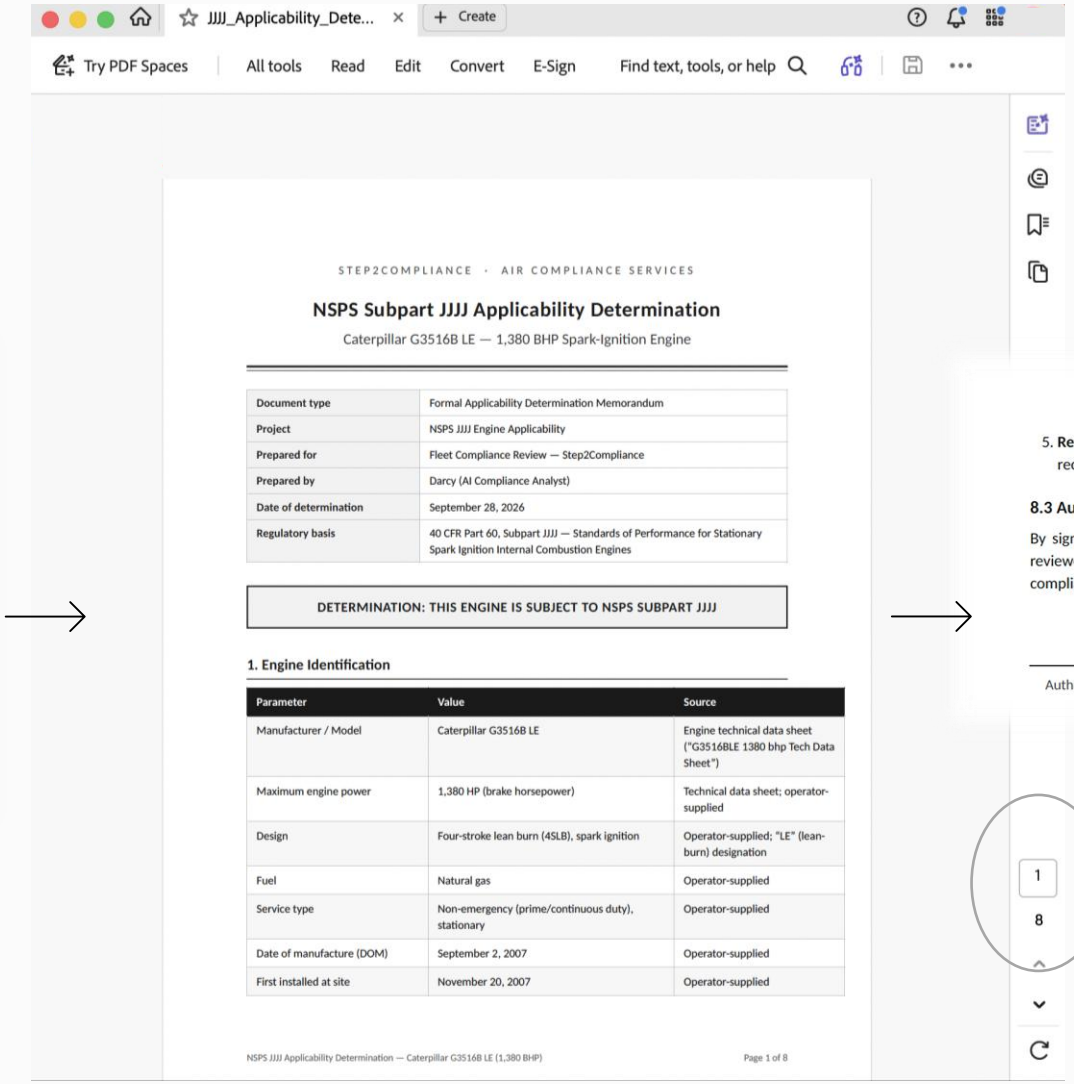
step2compliance

Agents

NSPS JJJJ Engine Ap...

NSPS JJJJ Engine Applicability

I need to determine NSPS JJJJ applicability for an engine in my fleet. It is a Caterpillar 1,380 HP unit with date of manufacture 09/02/2007. The engine was first installed on a site on 11/20/2007. The unit is non-emergency, natural gas fired, 4SLB. You have access to the JJJJ rule, rule history and summary, EPA implementation tools for the rule, and an engine technical data sheet. Do not go outside the boundaries and corpus defined within this project to find your answer. You may also use the step2compliance rule database to inform your answer. Take your time and perform a complete analysis. Is there anything else I should take into consideration while performing this analysis?



STEP2COMPLIANCE · AIR COMPLIANCE SERVICES

NSPS Subpart JJJJ Applicability Determination

Caterpillar G3516B LE — 1,380 BHP Spark-Ignition Engine

Document type	Formal Applicability Determination Memorandum
Project	NSPS JJJJ Engine Applicability
Prepared for	Fleet Compliance Review — Step2Compliance
Prepared by	Darcy (AI Compliance Analyst)
Date of determination	September 28, 2026
Regulatory basis	40 CFR Part 60, Subpart JJJJ — Standards of Performance for Stationary Spark Ignition Internal Combustion Engines

DETERMINATION: THIS ENGINE IS SUBJECT TO NSPS SUBPART JJJJ

1. Engine Identification

Parameter	Value	Source
Manufacturer / Model	Caterpillar G3516B LE	Engine technical data sheet ("G3516BLE 1380 bhp Tech Data Sheet")
Maximum engine power	1,380 HP (brake horsepower)	Technical data sheet; operator-supplied
Design	Four-stroke lean burn (4SLB), spark ignition	Operator-supplied; "LE" (lean-burn) designation
Fuel	Natural gas	Operator-supplied
Service type	Non-emergency (prime/continuous duty), stationary	Operator-supplied
Date of manufacture (DOM)	September 2, 2007	Operator-supplied
First installed at site	November 20, 2007	Operator-supplied

NSPS JJJJ Applicability Determination — Caterpillar G3516B LE (1,380 BHP) Page 1 of 8

5. **Retain this signed memorandum** with the verification records as the applicability determination of record for this engine.

8.3 Authorization

By signing below, the authorizer confirms that the verification actions in Sections 8.1 and 8.2 have been reviewed and completed as appropriate, and that this applicability determination is approved for use in the compliance record for this engine.

Authorizer signature | Date | Title

1
8

The accountable prompt

- States the asset: make, model, horsepower, manufacture date, site status
- Supplies the applicable rule text and internal records as source
- Asks for the citation and the reasoning, not just a conclusion
- Asks what additional information would change the answer
- Question, sources, answer, and check are all captured

It isn't the tool that got smarter; it's the human input behind it, and the trace that input leaves behind.

Demonstration two: an agent in a workflow

1

REQUEST & RETRIEVE

A person asks for a scheduled testing notification; the agent pulls the fleet record, requirement, and schedule from permitted systems.

2

DRAFT

It assembles the notification and shows exactly which records and rule text it used.

3

STOP

The action that would write to the record halts. The agent cannot proceed on its own authority.

4

APPROVE

A named, qualified human reviews the source, the draft, and the reasoning — then approves or rejects.

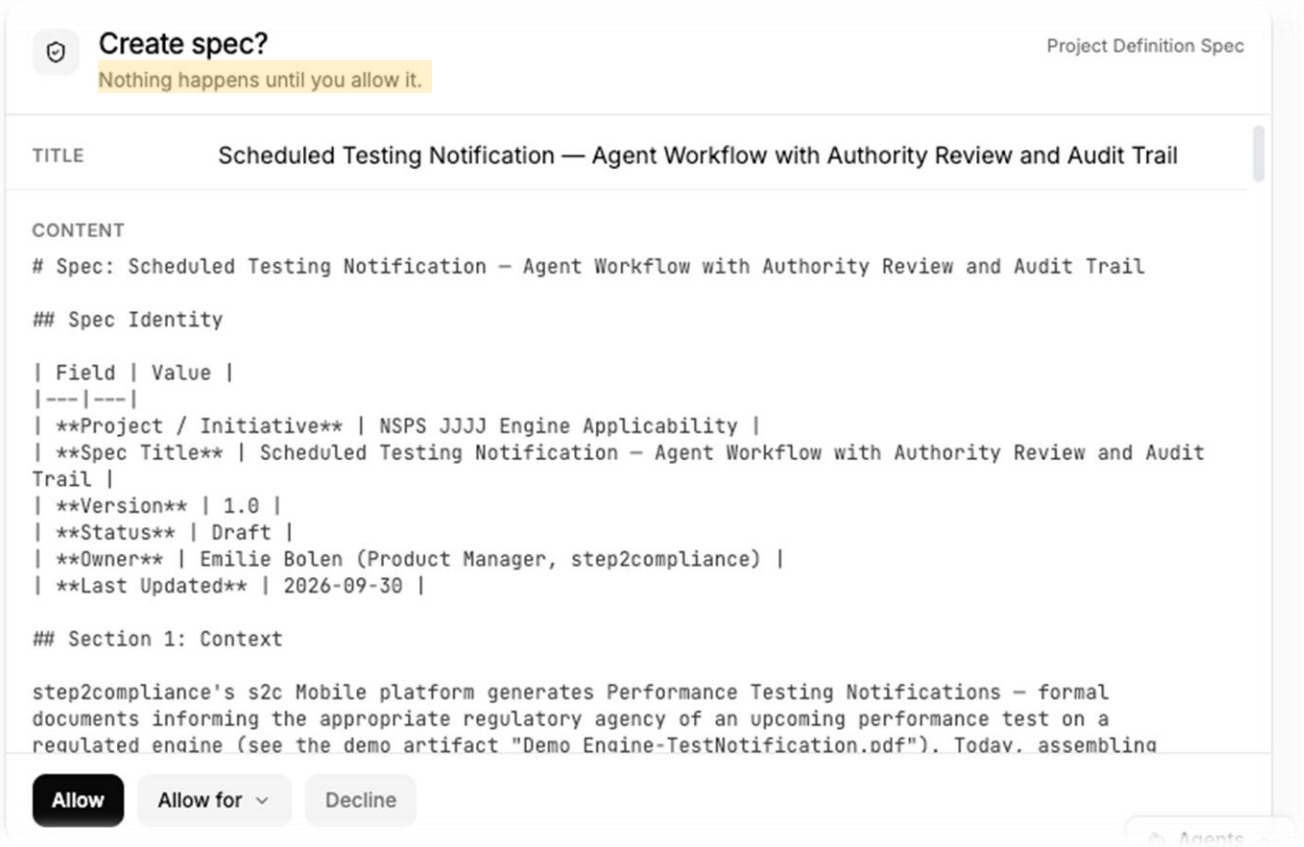
5

LOG

Request, sources, draft, approver, and timestamp are written to an auditable trail.

The two dark steps are the whole point. Remove them and you have automation. Keep them and you have accountable automation.

step2compliance



1

REQUEST & RETRIEVE

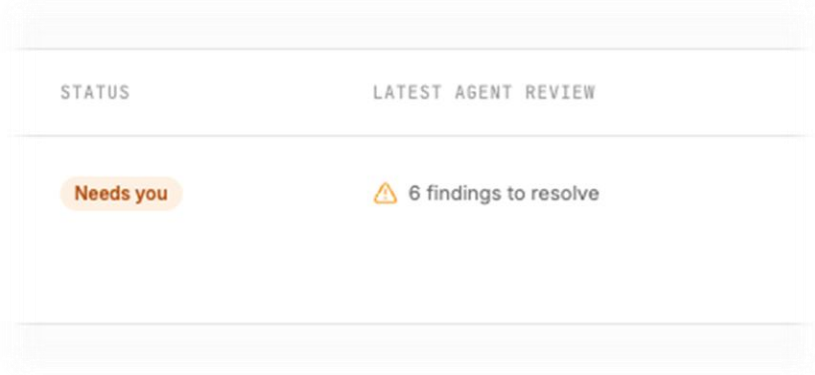
A person asks for a scheduled testing notification; the agent pulls the fleet record, requirement, and schedule from permitted systems.

2

DRAFT

It assembles the notification and shows exactly which records and rule text it used.

Automate the portion the agent can control with boundaries.



3

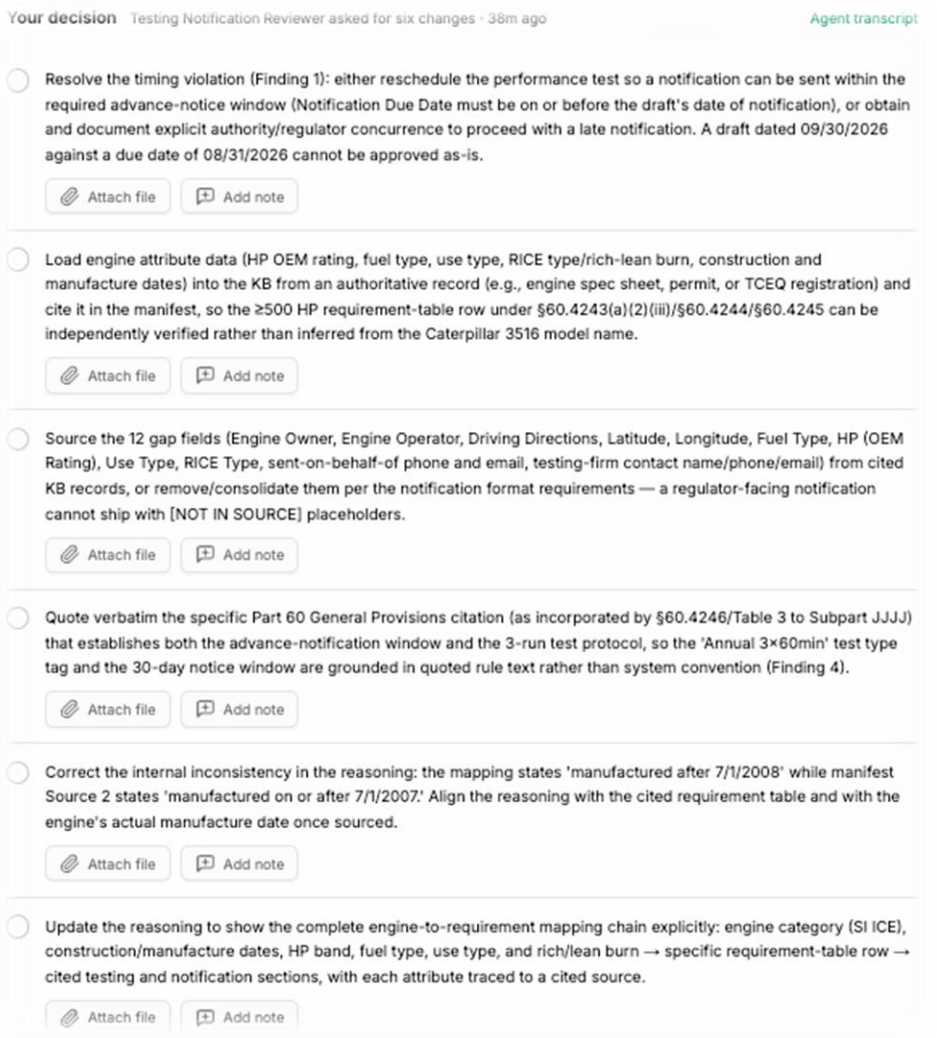
STOP

The action that would write to the record halts. The agent cannot proceed on its own authority.

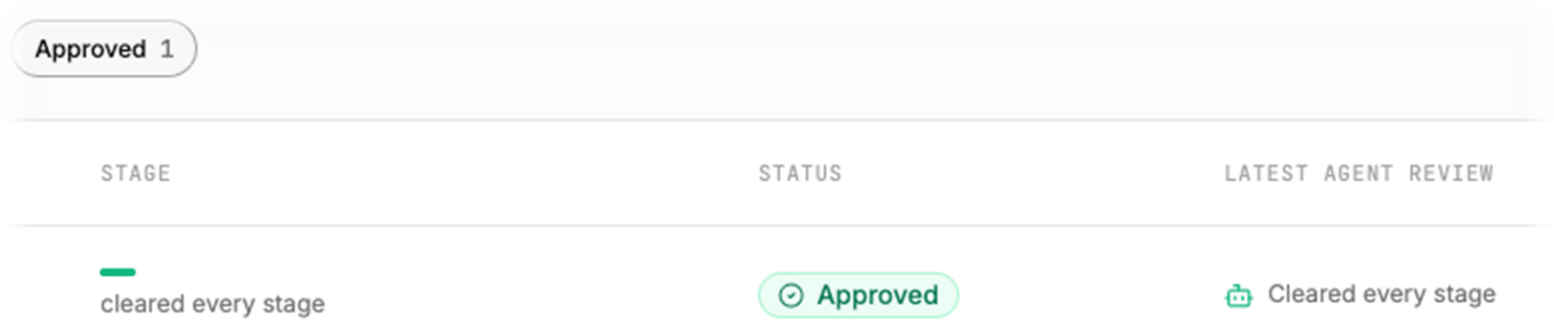
4

APPROVE

A named, qualified human reviews the source, the draft, and the reasoning — then approves or rejects.



The human gate in practice. Nothing is approved until you sign on it.



5

LOG

Request, sources, draft, approver, and timestamp are written to an auditable trail.



Performance Testing Notification

Performance Testing Notification
Engines subject to Testing & Notification Requirements of:

Federal - NSPS JJJJ Annual 3x60min

This is a notification to conduct a performance test on the engine and facility below scheduled for the following date:

Performance Test Date: 11/30/2026

Comments:

Date of Notification:	10/15/2026		
Engine Location:			
Facility Name:	Demo Facility	Latitude:	xx.xxxx
Engine Owner:	Owner 1	Longitude:	xx.xxxx
Engine Operator:	Operator 1	County:	Reeves
Driving Directions:		State:	TX
Engine Details:			
Unit ID:	Demo Engine	Make	Caterpillar
Serial #:	xxxxxxx	Model:	3516
Fuel Type:	NaturalGas	HP (OEM Rating):	1340
Use Type:	NonEmergency	RICE Type:	4SLB
Sent on Behalf Of:			
Company:	step2compliance		
Contact Name:	Emilie Bolen	Phone:	(866) 842-0660
Email:	emilie@step2compliance.com		
Testing Firm Information:			
Testing Company:	step2compliance		
Contact Name:	Emilie Bolen	Phone:	(866) 842-0660
Email:	emilie@step2compliance.com		

step2compliance © 2026 step2compliance, LLC (866) 842-0660 | info@step2compliance.com

AN AUDIT TRAIL YOU CAN TRUST

All steps documented, cited, and traceable. Agents that don't change the work you're doing but put framework around your process for an output that's the same *every time*.

Debrief: what accountability looked like

- A person decided in advance what the agent was allowed to do — and what it was not
- The question carried the context. The human supplied the asset detail, the rule text, and the records
- The tool was asked to show its sources and reasoning, and a qualified human checked the citation against the actual rule
- The action that would create a record stopped and waited for a named approver
- The exchange was logged, so the work can be explained months from now to someone who was not in the room

Not one of those behaviors was performed by the AI. The technology made the work faster; the people made it defensible.

Would you let what you just saw draft your next submittal?

Hands up for draft. Now keep them up only if you would also let it submit.

PART FIVE · 1:14 – 1:23

Frameworks and the Gap

Where accountability lives in a written framework, how fast both rulebooks are moving, and what your compliance stack has to do about it.

Brian Kromer with Emilie Bolen

Govern · Map · Measure · Manage

NIST gives accountability a place to live — voluntary, sector-agnostic, and free to cite.

1

GOVERN

Name the owner. This function explicitly assigns human accountability for AI risk — the federal government's own definition of IA.

2

MAP

Locate the risk. Inventory where AI is in use, on what data, and feeding which compliance obligation.

3

MEASURE

Test the output. Evaluate before you trust and keep evaluating after a model or vendor update.

4

MANAGE

Act on evidence. Treat what you find, set stopping thresholds, and define the incident path.

Minimum first step: assign a named human owner to every AI-assisted output that feeds a compliance obligation.

Source: NIST AI 100-1, AI Risk Management Framework 1.0 (Jan. 2023); NIST AI 600-1, Generative AI Profile (July 26, 2024) — names hallucination and confabulation among risks to manage

Both rulebooks are moving

- April 2026: EPA finalized amendments to NSPS OOOOb and EG OOOOc, revising flare and vent-gas net heating value provisions in response to reconsideration petitions
- Several OOOOb and OOOOc compliance dates were extended into 2026 and 2027, while every provision not extended stays in force — leak detection, monitoring and recordkeeping continue
- GHGRP reporting for calendar year 2025 was extended, with a proposed reconsideration of Subpart W pending and not finalized
- March 14, 2025: Congress used the Congressional Review Act to repeal the rule implementing the methane Waste Emissions Charge (further July action delayed the WEC until 2034)
- The AI rulebook is moving too: EU AI Act high-risk timing is being reset through the Digital Omnibus, and NIST has begun work on a trustworthy-AI profile for critical infrastructure

THE TRAP

A model trained before these changes will answer from the old rulebook — fluently, with a citation, and without telling you it is out of date.

Rule volatility is exactly where hallucination costs the most.

Source: EPA final rule, April 2026 (90 FR, reconsideration of OOOOb/OOOOc); EPA interim rule extending compliance dates; CRA joint resolution on the WEC implementing rule, Pub. L. 119-2, March 14, 2025; EU Digital Omnibus political agreement, May 2026

What a compliance system of record must do

AI is safe to use inside a system that already knows the rule, the owner, and the record. These five are the control layer.

1 CONTINUOUS RULE TRACKING

Someone must watch every jurisdiction and rule that touches your fleet, continuously. Our own platform tracks 130 jurisdictions and more than 1,200 existing and proposed rules.

2 APPLICABILITY WITH CITATION

The answer must arrive with the governing provision attached — NSPS JJJJ and IIII, RICE NESHAP ZZZZ, state rules — not a confident paraphrase.

3 WORKFLOW WITH A NAMED OWNER

Testing, inspections, maintenance, notifications and reporting routed to a person, with a due date and an approver.

4 CENTRALIZED, AUDITABLE RECORDS

Records commonly retained five years, organized in one place and producible on demand rather than reconstructed under audit pressure.

5 SECURE, LOGGED ACCESS

Role-based access for staff, vendors and clients, with a log of who did what and when. This is the evidence layer AI use now demands.

Disclosure: s2c Mobile is our platform, and the figures above are ours. The point is the category, not the vendor — whatever you use has to do these five things.

Five questions for any AI vendor

1 WHAT DID YOU MEASURE?

On whose data, over what period, and who audited the result? A number in a brochure is a marketing claim until someone shows the methodology.

2 WHERE DOES OUR DATA GO?

Is it used to train models, who can see it, how long is it retained, and can we get it deleted? Get the answer in the contract, not the demo.

3 WHAT HAPPENS WHEN IT DOES NOT KNOW?

Does the system cite its source, flag uncertainty, and abstain — or does it always produce a confident answer?

4 CAN YOU PRODUCE THE LOG?

Who asked what, what was returned, who approved it, and when. If the system cannot show that, it cannot support a certification.

5 WHO UPDATES IT WHEN THE RULE CHANGES?

How fast, by what process, and how do we know it happened? In our industry this is the question that separates tools from systems.

Usage is outrunning governance

Different populations, different question wording — do not combine them. The shared signal is acceleration.

21%

OF U.S. WORKERS

Say at least some of their work is done with AI, up from 16% a year earlier (Pew, Oct. 2025, n=5,010)

45%

USE AI AT WORK

At least a few times a year; 23% a few times a week or more; 10% daily (Gallup, Q3 2025, n=23,068)

~½

OF U.S. ADULTS

Now use AI chatbots, up from a third in 2024, with roughly one in four using them daily (Pew, June 2026)

~⅔

SAY IT IS TOO FAST

Roughly two-thirds of Americans say AI is advancing too quickly (Pew, June 2026)

Adoption accelerates, governance lags, and the public already knows it. Your organization is somewhere in that gap.

The consulting-firm governance statistics circulating online failed verification and are deliberately absent from this deck.

Source: Pew Research Center, Oct. 6, 2025 and June 17, 2026; Gallup, "AI Use at Work Rises," Q3 2025

Where AI belongs in a compliance workflow

Intake	AI assists	Organize incoming data, flag gaps and anomalies, summarize field notes
Research	AI assists	Locate applicable rule language, explain it plainly, compare requirements
Draft	AI assists	Assemble the first version of a report, notification, or checklist
Verify	Human only	Check every citation, number, and applicability call against the source
Approve & submit	Human owns	Acceptance, certification, submittal, and the five-year record — with AI use documented

Speed belongs in the first three steps. Judgment owns the last two. Confusing the two is how organizations get hurt.

PART SIX · 1:23 – 1:30

Practicing IA

Five practices, each one backed by a study you have already seen today — and a signature that stays yours.

All three presenters

Five practices you can take back Monday

1 CALIBRATE TRUST

Trust should track demonstrated reliability on your task type, not general AI capability. Confidence in the tool displaces checking; confidence in yourself restores it. (Lee et al., 2025)

2 MAKE CHECKING CHEAP

Overreliance is a cost-benefit choice. When explanations were easy to engage with, overreliance fell to zero in the authors' exploratory analysis. Put the raw data one click from the conclusion. (Vasconcelos et al., 2023)

3 REWARD ACCURACY

If your bonus system punishes people for overriding the algorithm and being right, you have trained the organization to stop thinking. (ESMT / Frankfurt School, 2024)

4 NAME THE ACCOUNTABLE USER

The JAMIA review lists emphasizing user accountability as a demonstrated mitigator of automation bias, alongside training and design. (Goddard et al., 2012)

5 NEVER LET IT SELF-POLICE

ChatGPT recognized its own hallucinations only 62.59% of the time in QA. Verification must be external — the source data, the regulation, the instrument reading. (HaluEval, 2023)

Verify before you sign

- [ ] Every citation opened and read in the actual rule, permit, or spec — not accepted because it looked right
- [ ] Every number traced to a source record, and every calculation reproduced independently
- [ ] Dates and versions confirmed current; ask the tool directly whether the rule has been amended, then check anyway
- [ ] Scope confirmed and gaps resolved: right unit, right site, right period — with no assumption quietly filled in by the model
- [ ] The check itself recorded — what was verified, by whom, against what, on what date

Verification is not distrust of the tool. It is the price of putting your name on the output.

TAKE THIS BACK WITH YOU

Your personal AI charter

I will declare it.

I will say where AI helped, in any work that carries my name.

I will verify it.

I will check every citation, number, and rule reference against the source.

I will contain it.

Client and company data goes only into approved tools.

I will document it.

I will record what I checked, so the work can be defended later.

I will own it.

If it is wrong, it is mine. The AI did it is not available to me.

Five sentences. Sign them for yourself before you ask anyone else to.

Your signature, your world

1

THE PERSON WHO SIGNS

The emissions report that an AI helped assemble still carries one human name on the certification line.

2

THE TECHNICIAN AT 2 A.M.

An AI-generated alert fires. Someone decides whether to roll a truck, shut a unit, or wait for daylight.

3

THE ENGINEER ON THE PACKAGE

An AI maintenance recommendation arrives on a compressor package. Accepting it is a decision, not a click.

The AI can draft.
Only you can attest.

THE CLOSE

Refuse the deal.

Science says that using a sycophantic AI makes people less willing to take responsibility — and more certain they are right.

That is the deal on offer: easier work, weaker accountability.

Your job — and your signature — is to refuse that deal.

Every figure in this course was verified against its primary source. The full source pack, with links, is yours.

Got Questions?

Brian Kromer

Founder & Managing Director
step2compliance

bkromer@step2compliance.com
970.227.5935

Emilie Bolen

Compliance Services Manager
step2compliance

emilie@step2compliance.com
970.714.4234

Jeremiah Cantu

Senior Technical Architect
Innovative Solutions, AWS Premier Partner

jcantu@innovativesol.com
830.237.9057

step2compliance · 4818 Snowdrift Circle, Fort Collins, CO 80528 · 866.842.0660

WWW.STEP2COMPLIANCE.COM

Primary-verified source pack available — every citation, every link, with verification status noted.

Take it with you · step2compliance.com/ai

APPENDIX · NOT PRESENTED

Take-Home Toolkit

The material below is for your copy of the deck: the tools to run this at home, the full source index, and the numbers we deliberately left out.

Reference only — skip in the room

Private AI and public tools: what changes

Dimension	PRIVATE, GOVERNED DEPLOYMENT	CONSUMER TOOL ON A PERSONAL LOGIN
Where your data goes	Into a controlled environment you define	Into a vendor platform under consumer terms
Training on your input	Excluded by configuration and contract	Varies by tier and setting; often opted in by default
Access control	Role-based, provisioned, revocable	Whoever has the password, including former employees
Auditability	Session logs you can produce on request	Typically, none you can hand to an auditor
Cost of getting it wrong	Contained by design	Client data outside your control, permanently

If you have not chosen, your people have already chosen for you.

Five clauses to write this quarter

1	APPROVED TOOLS AND DATA	A named list of what may be used for company work, and what may never be entered into any AI tool: client data, personnel data, anything privileged.
2	DISCLOSURE	When and how employees declare AI assistance in work products and records.
3	VERIFICATION AND APPROVAL	What must be checked before AI-assisted output becomes a deliverable, and which rungs of the ladder require a named approver.
4	INCIDENT REPORTING	How to report an AI error or data exposure without fear — fast reporting beats quiet cleanup.
5	REVIEW CADENCE	Who revisits this policy, how often, and what triggers an off-cycle review.

A one-page policy your people read beats a forty-page policy your counsel admires.

Monday morning: a 30-60-90 plan

FIRST 30 DAYS

- Ask your team, without penalty, what AI tools they already use
- Name an executive sponsor and an AI steward
- Publish an approved-tool list and a one-page data rule
- Run the vocabulary and the ladder past the whole team

BY DAY 60

- Write the five policy clauses; keep it to one page
- Pick one workflow and mark where AI assists and where a human signs
- Stand up a governed environment for anything touching client data
- Name the verifiers for each discipline

BY DAY 90

- Log real usage and review the errors and near misses captured
- Train the whole company, not just technical staff
- Add AI-use documentation to your record practices
- Set the review cadence and calendar the next review

None of this requires a new budget line. All of it requires somebody to own it by name.

Source index — the hook and the risk

EPA Oil & Gas Alternative Test Methods
The MATM list, MATM-001 to MATM-018 — epa.gov/emc/oil-and-gas-alternative-test-methods

EPA Clean Connect approval letter
LeakFinder as alternative test method, 40 CFR 60.5398b(d), May 13, 2025

EPA Subpart W / GHGRP final rule
Empirical data, advanced technologies, super-emitter program, May 6, 2024

PHMSA R&D projects 953 and 954
AI-enabled pipeline risk modeling; KAI-MAP hydrogen integrity — primis.phmsa.dot.gov

GTEN 2023, paper 23-GTEN-303
Hybrid AI/physics turbine model, tested in TC Energy Canada Gas Compression

Williams case study, June 17, 2026
Field checklist and software delivery figures — company-published, not independent

Lee et al., CHI 2025
Generative AI and critical thinking; 319 workers, 936 examples — DOI 10.1145/3706598.3713778

Cheng et al., Science 2026
Sycophantic AI and personal responsibility — DOI 10.1126/science.aec8352

Goddard, Roudsari & Wyatt, JAMIA 2012
Automation bias systematic review, 74 studies — DOI 10.1136/amiajnl-2011-000089

Vasconcelos et al., CSCW 2023
Explanations reduce overreliance; N=731; zero overreliance when checking was easy

Qazi et al., 2025
Automation bias in AI-trained physicians — DOI 10.1101/2025.08.23.25334280

Magesh et al., JELS 22:216–242 (2025)
Legal AI tools hallucinate 17–33%; Dahl et al. 58–82% for general models

Source index — consequences, frameworks, industry

Mata v. Avianca, 22-cv-01461 (S.D.N.Y.)

Rule 11 sanctions opinion, June 22, 2023 — paraphrase only, opinion not machine-read

Gauthier v. Goodyear, E.D. Tex.

Memorandum and Order, November 2024 — $2,000 penalty and mandatory AI CLE

S.D.N.Y. Judge Cronan model certification

AI use in filings, revised October 23, 2025 — nysd.uscourts.gov

CIT Judge Kelly, Order on AI

Tool disclosure, AI-drafted text, confidentiality certification, July 2025

NIST AI 100-1

AI Risk Management Framework 1.0, January 2023 — Govern, Map, Measure, Manage

NIST AI 600-1

Generative AI Profile, July 26, 2024 — names hallucination and confabulation

Pew Research Center

Oct. 6, 2025 (21% of workers, n=5,010); June 17, 2026 (about half of adults)

Gallup, Q3 2025

AI use at work: 45% at least occasionally, 23% weekly, 10% daily; n=23,068

Vectara Hallucination Leaderboard

Document-summary error rates, 1.8–5.1% for leading models

OpenAI GPT-5 System Card

SimpleQA without browsing: gpt-5-thinking accuracy 0.55 (55%), hallucination 0.40 (40%)

EPA final rule, April 2026

Reconsideration of OOOOb/OOOOc — flare and vent-gas net heating value provisions

Market reporting on AI gas demand

RBC Capital Markets 2026; GE Vernova Q1 2026; S&P Global — analyst material, not primary

Numbers we deliberately excluded

Each of these circulates widely and could not be confirmed in a primary source. None appears anywhere in this deck.

"90% of workers use AI weekly, 50% daily"	Attributed to Pew. Pew's actual finding is 21% of workers; Gallup reports 45% occasionally, 10% daily.
"–27% critical thinking, 2.3× acceptance"	Attributed to Lee et al. Not in the paper. The finding is directional and is presented that way.
Consulting-firm governance statistics	"84% want training," "17% of C-suites have governance" and similar — unverifiable; dropped entirely.
"12% lower bonuses," "120 workers," "70%"	Attributed to the ESMT study. Not in the press release. The finding is cited qualitatively instead.
Vendor ROI claims for AI platforms	"72% less downtime," "340% ROI" and similar — published by the sellers, not audited. Category cited, number omitted.

A deck about accountability has to be accountable for its own numbers.

Where to verify, yourself

1 EPA TEST METHODS

epa.gov/emc/oil-and-gas-alternative-test-methods — the live MATM list, with approval dates and thresholds.

2 EPA OIL AND GAS RULES

epa.gov/controlling-air-pollution-oil-and-natural-gas-operations — OOOOb/OOOOc status, amendments and extended dates.

3 NIST AI FRAMEWORKS

nvlpubs.nist.gov — AI 100-1 and AI 600-1. There is no verified AI RMF 1.1; ask any vendor citing one for the document.

4 FEDERAL COURT ORDERS

nysd.uscourts.gov and cit.uscourts.gov — the AI certification forms, in their original text.

5 PHMSA RESEARCH

primis.phmsa.dot.gov/rd — project status for AI-enabled risk modeling and leak detection.

Every claim in this course points to a document you can open without our help. That is the standard we are asking you to hold your tools to.