ONE-DAY TRAINING PROGRAMME · 8 HOURS · ONSITE OR LIVE VIRTUAL

Validating AI in GxP

Put AI and machine learning into GxP with evidence that holds up under inspection. One live case runs through the whole method, and your team leaves able to work it themselves.

300+
regulatory inspections supported, with zero critical findings

Annex 22
contributor to the EU GMP industry review of the AI annex

Author
Validating AI in GxP: A Practitioner's Guide

WHY NOW

The rulebook arrived in the last 18 months

Draft Annex 22, final FDA CSA guidance and the FDA and EMA joint AI principles all landed since 2025. In April 2026 the FDA issued its first warning letter citing AI use in drug manufacturing, after a firm let an AI agent write its GxP documents and took the agent's word for what the regulations required. Waiting for final guidance now means carrying the risk. Teams are deploying AI in quality operations either way, so the untrained team is the exposure, not the tool.

WHO IT IS FOR

The people who will have to defend it

Validation and CSV leads, QA and quality heads, IT compliance and data integrity owners, Qualified and Responsible Persons, data scientists working in regulated contexts, and the reviewers who will stand behind an AI system in an inspection. Bring the decision makers and the people who build the evidence together. Up to 30 people in one sitting.

THE PROBLEM WE SOLVE

A trained CSV instinct fails on AI, for three specific reasons

It changes after release
A system that keeps learning after sign off is one you validated once and no longer fully know.

It is opaque
You can read every line of traditional software. You cannot do that with a model.

Its behaviour is inherited
Behaviour comes from the data the model learned on, not only from the code it runs.

The April 2026 warning letter is narrower and more useful than the headlines made it. FDA's position, in its own words: any output or recommendation from an AI agent must be reviewed and cleared by an authorised human representative of the firm's quality unit. Human accountability at the quality unit is exactly what this programme builds. (Warning letter 320-26-58, 2 April 2026.)

WHAT WE TEACH

TRUST: five decisions your team makes on the job

T	R	U	S	T
Tier the risk	**Requirements & intended use**	**Understand the data**	**Substantiate with evidence**	**Track in operation**
"What tier is this, and who signs?"	*"What am I claiming, and where does it stop?"*	*"Can I trust the data it learned from?"*	*"What evidence earns the validated state?"*	*"How do I keep it inside the envelope?"*

Governance decides whether the system may exist. TRUST earns and keeps its validated state. The double T is deliberate: the sequence opens and closes on "can you trust this system", the question every inspector is really asking. One live case runs through every letter, a QMS deviation-to-action agent, and we plant the flaws on purpose: the representativeness gap, the drift excursion, the unsupported CAPA text and the missing independence signature.

THE AGENDA

Every session opens on a Monday morning decision, not a topic

THE PROGRAMME		09:00 to 17:00 · 8 hours · onsite or live virtual · up to 30 people
09:00	**Why does my CSV instinct fail here?**	The three breaks. What is new, and what is not.
09:45	**May this system even exist?**	Governance before validation. The operating model and who owns it.
10:30	**What tier, and who signs?**	The four-tier risk model. Delegates classify the case and defend it.
11:30	**What am I claiming, and where does it stop?**	Intended use statements and the human-in-the-loop boundary.
12:15	**Where do Annex 22, CSA, Part 11 and the AI Act apply?**	Positioning the rulebook, not reciting it.
13:45	**Can I trust the data it learned from?**	Representativeness, provenance, labelling, test data independence.
14:45	**What evidence earns the validated state?**	The validation plan, acceptance criteria, the three independences.
15:45	**How do I keep it inside the envelope?**	Monitoring, drift thresholds, retraining triggers, audit trail.
16:30	**Am I inspection ready?**	The evidence pack, the question bank, and each delegate's commitment.

RECOMMENDED UPGRADE · ADD DAY TWO

Day one teaches the method. Day two is where the behaviour changes. Teams run TRUST on their own live AI use case with coaching, work a data and evidence clinic on their own datasets and gaps, sit a mock inspection beat drawn from the 120-question bank, take on the agentic and LLM frontier, and leave with a written first-100-days plan and a named owner per action. Take it when an AI use case is already live, and whenever we deliver onsite, because a single day rarely justifies the trip for either side.

WHAT EACH PERSON WALKS OUT WITH

- Tier an AI system and say who signs
- Write intended use that bounds the claim and the effort
- Judge data quality, lineage and test data independence
- Build the AI validation plan, acceptance criteria and evidence pack
- Set up monitoring, drift response and the AI risk register
- Answer an inspector, and leave with a 30-day commitment on a named system

WHAT IS INCLUDED

- Live delivery by Sachin Bhandari, onsite or as a virtual cohort
- The TRUST toolkit: classification, intended use, evidence and monitoring templates
- The golden case walkthrough, or your own case on request
- Recording for virtual cohorts, and a certificate of participation
- A pre and post competency check
- A follow-up question and answer window

WHAT THE ORGANISATION GETS, STATED PLAINLY

AI projects stop stalling at the quality gate

Teams who can tier and bound a system approve it in weeks, not months.

An inspection answer that already exists

Evidence pack, risk register and monitoring record, built the way an inspector reads them.

One vocabulary across quality, IT and data science

Most AI governance failures are handover failures, not technical ones.

Less dependence on outside help

The capability stays in the building, so the next AI system starts without a consultant.

DELIVERY AND LOGISTICS

What booking one actually involves

Format

Onsite at your site, or delivered as a live virtual cohort. Virtual widens reach across sites; onsite adds the energy and the tailoring of a room.

Group size

Up to 30 people in one sitting. Further sittings are arranged separately.

Lead time

A tailored proposal within one working day. Dates are held on a released purchase order or advance, and most sessions deliver within four to six weeks of booking.

Tailoring

Built on yours. Sessions work on your system inventory, your SOPs and your inspection history, and teams leave with output they can use the next working day.

QUESTIONS WE ARE USUALLY ASKED

Is it generic, or built on our systems?

Built on yours. Sessions work on your system inventory, your SOPs and your inspection history, and teams leave with output they can use the next working day.

How fast can a date be confirmed?

You receive a tailored proposal within one working day. Dates are held on a released purchase order or advance, and most sessions deliver within four to six weeks of booking.

What do participants take away on paper?

The working templates for the framework taught, the worked case built in the room, and a certificate of participation. Virtual cohorts also receive the recording.

YOUR TRAINER

Sachin Bhandari

Founder and Principal Consultant, TrustBridge Compliance. Sachin has spent 25 years in pharma quality IT. He wrote *Validating AI in GxP: A Practitioner's Guide*, the field's first practitioner handbook, and contributed to the EU GMP Annex 22 (AI) industry review. He has validated 20+ AI applications in live GxP use, deviations, queries and CAPA support among them, and taken 9 of them through health-authority inspection, 6 FDA and 3 MHRA, with zero compliance observations. He teaches practitioners and briefs boards, often in the same room.

Next step. See all three in-house programmes, what each day covers and what a session costs, and ask for a date. A tailored proposal follows within one working day.

trustbridge-compliance.com/book-a-training