

Is AI ready for our industry?

One unsettled argument, and far less running than the noise suggests

An ARTO Report
August 2026

ARTO

Contents

The Episodes 03

What the industry does not talk about: the failures 04

Does AI need a data foundation first? 05

What is actually running 07

The one pilot that met a regulator 09

Where the regulator has drawn the line 10

Whether a human stays in the loop 12

Where to start 14

The Episodes

All forty episodes reviewed for this report, and the full set every quotation is drawn from. Quotations are not attributed to individuals. All the episodes are free to watch, and each thumbnail links to one.



What the industry does not talk about: the failures

An industry event director was asked what conversation the industry needs to have more openly. She did not name a technology. She named a bias.

> *AI and automation still really new and we hear stories of all the successes and what it is that's happening but we don't hear those what didn't work and where are you struggling.*
>
> **Industry event director**

Honest conversations get messy, she added, but talking about where we fail helps everybody learn faster.

The corpus proves her point. Across 40 conversations with CEOs, site leaders, quality heads and validation specialists, exactly one guest describes an AI programme that did not work as intended. Everything else is a success, an intention, or a general observation about the direction of travel. That imbalance is worth noticing, because the successes are narrower than the enthusiasm around them, and the single failure is more useful than any of them.

Does AI need a data foundation first?

The real division is not optimism against pessimism. Nearly everyone expects AI to matter. The division is about sequence.

A principal consultant who has spent a career in IT and regulatory compliance in life sciences, and who sits on an industry standards committee, puts the case that data comes first, without hedging.

> *I think there's one real important driving technology which everyone's talking about which is AI and AI is massively dependent upon huge volumes of very high quality data.*
> **Principal consultant, IT and regulatory compliance**

> *Without those huge volumes of very high quality data, AI is not going to achieve anything like the benefits it could.*
> **Principal consultant, IT and regulatory compliance**

He is not a sceptic. Asked whether AI would be game changing in manufacturing, he said yes, it would, and called it a no-brainer, adding that quite how it changes the game was not yet clear to anyone. His point is about what the technology needs in order to work.

> *You can't just throw AI at a problem and it will solve it for you.*
> **Principal consultant, IT and regulatory compliance**

To make it work, he says, you need a solid foundation of data that is high quality, well structured and properly understood.

A project manager at a large pharmaceutical company disagrees, on the record, in a different episode. He has over 30 years in commissioning, qualification and validation, and he treats the foundations argument as a reason the industry gives itself for waiting.

> *If you're not using AI, you need to because five years from now, it's going to be all over.*
> **Project manager, large pharma**

Put to him that much of the industry is still on paper and is being asked to run before it can walk, he pushed back directly.

> *I've got a different approach to that because for me there's again there's two sides of AI.*
>
> **Project manager, large pharma**

> *Even if you're purely paper and you have no digital data, you can still say here's a PDF and an LLM will take PDFs showing all my requirements.*
>
> **Project manager, large pharma**

In twenty minutes, he says, he can produce all of this without any digitised data at all, and it moves his validation work forward.

Both are right. They are describing different problems. The first means analytics over plant and process data, where a model is only as good as the record it learns from. The second means document generation, where the inputs are a requirements specification, an equipment manual and a template, and where nothing needs to be digitised first because the source material is already a document.

Sort candidate use cases into those two buckets before sorting them by value. One bucket has a prerequisite. The other does not.

What is actually running

Separating deployment from intention thins the picture considerably. Six organisations describe AI genuinely in use.

A contract research organisation has a project classifying full, empty and intermediate viral vectors from cryogenic electron microscopy imagery. Its quality and compliance lead presents it as beyond human capability rather than merely faster.

> *We have a project... to identify full and empty vectors and also identify the intermediates because a human cannot do that because you have to analyze thousands of pictures and the differences are really just a shade of gray.*
>
> **Quality and compliance lead, contract research**

> *There is an expert who really said no, without AI this would not be possible.*
>
> **Quality and compliance lead, contract research**

The clearest commercial case comes from a CDMO chief executive, who describes a partnership using an algorithm for formulation development, and claims a repeated result rather than a pilot.

> *We've demonstrated on many program that it works and we are able to shrink development of the formulation time.*
>
> **Chief executive, CDMO**

He is careful about what it has changed commercially. Asked whether customer expectations had shifted with AI, he said no, they were still the same, and customers still expect performance and excellence in delivery.

Another chief executive describes AI screening materials and suppliers for his company's sustainability strategy, replacing traditional life cycle assessments and supplier surveys, plus an electronic lab notebook with AI functionality used for data mining.

> *We've been pretty actively using AI from our sustainability strategy actually in collaboration with a company...*
>
> **Chief executive, CDMO**

The project manager describes his employer running an internal AI tool inside the firewall, and his own use in detail: generating qualification protocols from requirements and equipment manuals, drafting standard operating procedures, translating them, and building training programmes with quizzes. A manufacturing technology manager at a biopharma company reports report writing, protocol writing and automated statistics, and is explicit that this is general use, not cleaning validation. A computerised systems validation specialist used AI to check how standards had shifted over 15 years while updating an industry good practice guide.

Everything else is planned or hoped for, and should be labelled as such. The chief executive of a biologics CDMO announced an AI version of his company's modular design framework during his episode.

> *This is first time I mentioned this actually in a broadcast but our version with AI is the next version that you would see very soon.*
>
> **Chief executive, biologics CDMO**

A division president at a CDMO makes the largest claim in the corpus about AI's significance.

> *The globalization of AI is happening faster than the internet happened faster than electricity.*
>
> **Division president, CDMO**

Asked what his company is doing, the answer is a start with an unnamed integration partner, a corporate rollout described as moving forward, and a hire under consideration.

> *We're considering even bringing in somebody to lead the AI journey for us.*
>
> **Division president, CDMO**

That gap between claim and deployment is not a criticism of any one company. It is the shape of the sector right now, and reading it accurately is the difference between a strategy and a press release.

The one pilot that met a regulator

A board member at a pharmaceutical manufacturer, after a career that included senior manufacturing leadership at a large pharmaceutical company, gives the only first-hand account in the corpus of an AI programme colliding with a regulator.

> *In my previous job... we started our first AI pilot and that was linked to our 100% visual inspection, so it's about picture recognition.*
>
> **Board member, pharmaceutical manufacturer**

> *In beginning we thought this is going to take over basic critical safeguard release of the product and then we came to learn by working with authorities and you will not be able to go there in any shape of form.*
>
> **Board member, pharmaceutical manufacturer**

The programme was not cancelled. It was repointed. Instead of deciding release, the AI engine gave engineers recommendations on how to adjust the masks that determine whether a product passes or fails, which addressed the problem they actually had: scrapping too much good product.

The lesson he draws is the sharpest sceptical argument in the corpus, and it is not about trust in the technology.

> *If it's not very, very, very strictly defined you cannot use it.*
>
> **Board member, pharmaceutical manufacturer**

He is not fearful that AI will take over, he says, because most processes are not yet disciplined enough for AI to be involved at all. Working with any modern technology, in his account, exposes the weaknesses in your own processes before it delivers anything.

That reframes the readiness question. The barrier is not the model. It is the process definition underneath it. He is not against the technology either: AI, he says, can help us fail faster so we do not waste time.

Where the regulator has drawn the line

On 2 April 2026 the FDA issued a cGMP warning letter to Purolea Cosmetics Lab containing, for the first time, a dedicated finding on inappropriate use of artificial intelligence, citing 21 CFR 211.22(c). AI agents had generated drug product specifications, standard operating procedures and master production and control records, and those documents entered use without review and approval by an authorised representative of the firm's quality unit.

The line is narrow and specific. It is not that AI may not be used. It is that AI output entering a quality system requires human review, documented approval and quality unit accountability.

One guest raised the letter unprompted, and his reading of it is worth keeping even though his recollection of the details differs from the record. His conclusion was that the failure was the absence of a documented process, not the absence of a person.

> *I mean, they were probably were reviewing it or something, but they didn't have a really good process for it, and that's what they got them on.*
>
> **Project manager, large pharma**

A quality and compliance lead places AI inside a broader shift already under way, from validating every button in a piece of software towards a risk based approach, driven by GAMP 5 second edition, critical thinking and computer software assurance. He also flags the cost of that freedom: with judgement replacing prescription, you are always at risk of not pushing all the buttons the regulators would expect.

A co-chair of an industry standards steering committee confirms that guidance is coming.

> *There is going to be a new one on AI and ML as a good practice guide.*
>
> **Industry standards committee co-chair**

He mentions in passing that the EMA presented at a recent conference on the Annex 11 draft revision and on the new Annex 22. That is the extent of it. Annex 22 is named and nothing about its content is discussed anywhere in the corpus.

A manufacturing technology manager describes the practical consequence.

> *Typically, technological advance is much faster than authorities follow, which is why we're still being conservative.*
>
> **Manufacturing technology manager, biopharma**

Whether a human stays in the loop

On whether a human stays in the decision loop, the corpus is close to unanimous. The project manager wants a documented human between output and action. The board member sees no problem with a person applying the final release stamp, while expecting AI to play a growing role up to that decision. A global quality engineering lead at a large pharmaceutical company is blunt.

> *In my opinion, still we're way far from relying 100% on AI and we should never rely 100% on AI.*
>
> **Global quality engineering lead, large pharma**

His concern is complacency rather than capability. Rely on it heavily, he warns, and people start taking results for granted.

A chief executive draws the boundary by task rather than by principle.

> *I think we also want to be very careful though because I don't believe in AI writing reports.*
>
> **Chief executive, CDMO**

Standard jobs that require no thinking and no communication will certainly be automated and replaced by AI, he expects. The rest he draws a line under.

> *Things like complexity management, problem solving, creativity, creative solutions will remain a human domain hopefully.*
>
> **Chief executive, CDMO**

The quality and compliance lead makes the same point from the technology's side.

> *Even if you have the large language models... what they are missing is creativity.*
>
> **Quality and compliance lead, contract research**

Large language models, in his account, take existing knowledge and existing data and rearrange it. Whether that counts as creativity, he says, he is not 100% sure.

There is exactly one dissenter, and it is a guest with real authority over how these systems will be validated. Discussing automated code review, a computerised systems validation specialist describes an end state without much human involvement in it at all.

> *When you've got 3,000 pages of code to review you get a bit weary, don't you?*
>
> **Computerised systems validation specialist**

> *AI just doesn't seem to get weary.*
>
> **Computerised systems validation specialist**

> *Once it's learned the rules it's looking for... ultimately I think we'll get to the situation where one organization's computers are just talking to some another organization's, there won't be that much human interaction.*
>
> **Computerised systems validation specialist**

That forecast deserves weight because of where it comes from. It also deserves the qualifier she supplies herself.

> *I haven't got examples of it being used but certainly conversations about how it was going to be used.*
>
> **Computerised systems validation specialist**

Where to start

The disagreement between the two practitioners resolves into a sequence rather than a choice.

Sort candidate use cases by whether they consume plant data or documents. Document generation, protocol drafting, code review, translation and training material need no digitisation programme first, and are available to a paper site this quarter. Anything that learns from process, batch or quality data inherits every weakness in that data, and the data prerequisite applies in full.

Then apply the board member's test. If a process is not defined tightly enough to describe unambiguously, no model will rescue it, and the attempt will surface the gap rather than close it.

Scope the human review before the tool, not after. The FDA's first AI warning letter turned on documented quality unit review of AI output entering the quality system. That control is cheap to design at the start and expensive to retrofit.

And expect the compliance function to be an ally rather than an obstacle, which is where the data-first argument ends up.

> *The champions of AI and us who have been knocking around in data compliance, regulatory compliance for some time, we're actually now finding ourselves pulling in the same direction.*
>
> **Principal consultant, IT and regulatory compliance**



ARTO

arto-talent.com