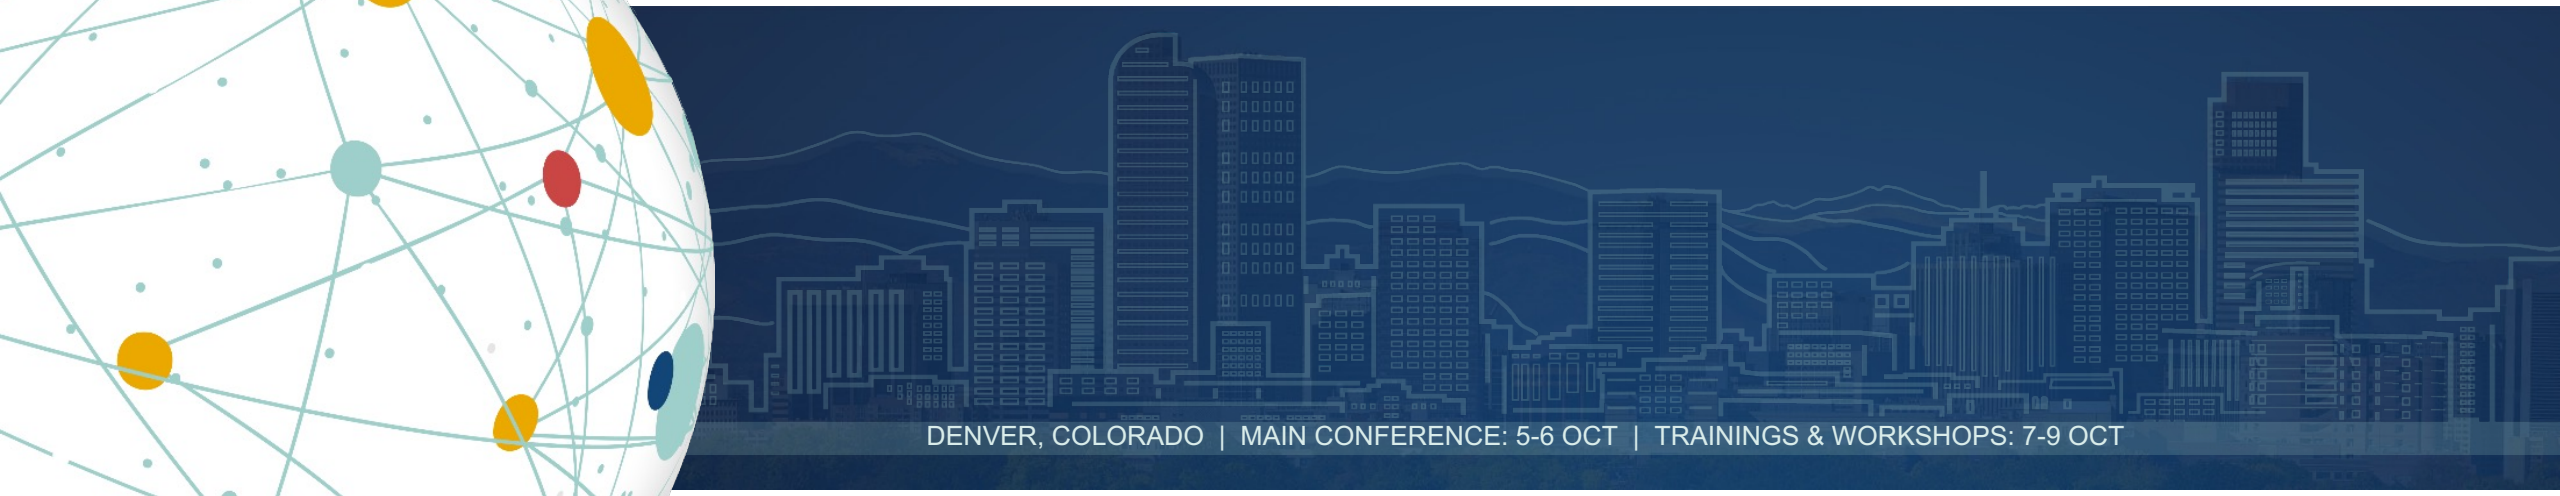


2026 US Interchange

THE FUTURE IS CONNECTED: STANDARDS AND AI POWERING DIGITAL TRANSFORMATION



Same Team, Different Worlds

Building AI competency across every level of the TMF function

Rob Jones, Arithmos Tech

ARITHMOS

Disclaimer

CDISC

The views and opinions in this presentation are the presenter's own and do not necessarily reflect the policy or position of CDISC. Mention of any product, vendor or service is not an endorsement by CDISC.

Arithmos Tech

This presentation reflects the presenter's experience and does not represent Arithmos Tech or its clients. Examples are illustrative. Nothing here is regulatory, legal or compliance advice.

About me



Rob Jones
Clinical Business Partner
Arithmos

- Took the red pill and is now one of those annoying AI guys “Have you heard about Claude?”
- Has multiple agents running, all with their own names and personalities
- Oh and also works very closely with the TMF and other Clinical Operation functions

Agenda

01 Where we are

The timeline
Why it matters to TMF teams

02 The levels

Five levels, one team
What happens when they collide

03 Building capability

Training, built from the foundation up
Role vs competency
Career progression: we have done this before

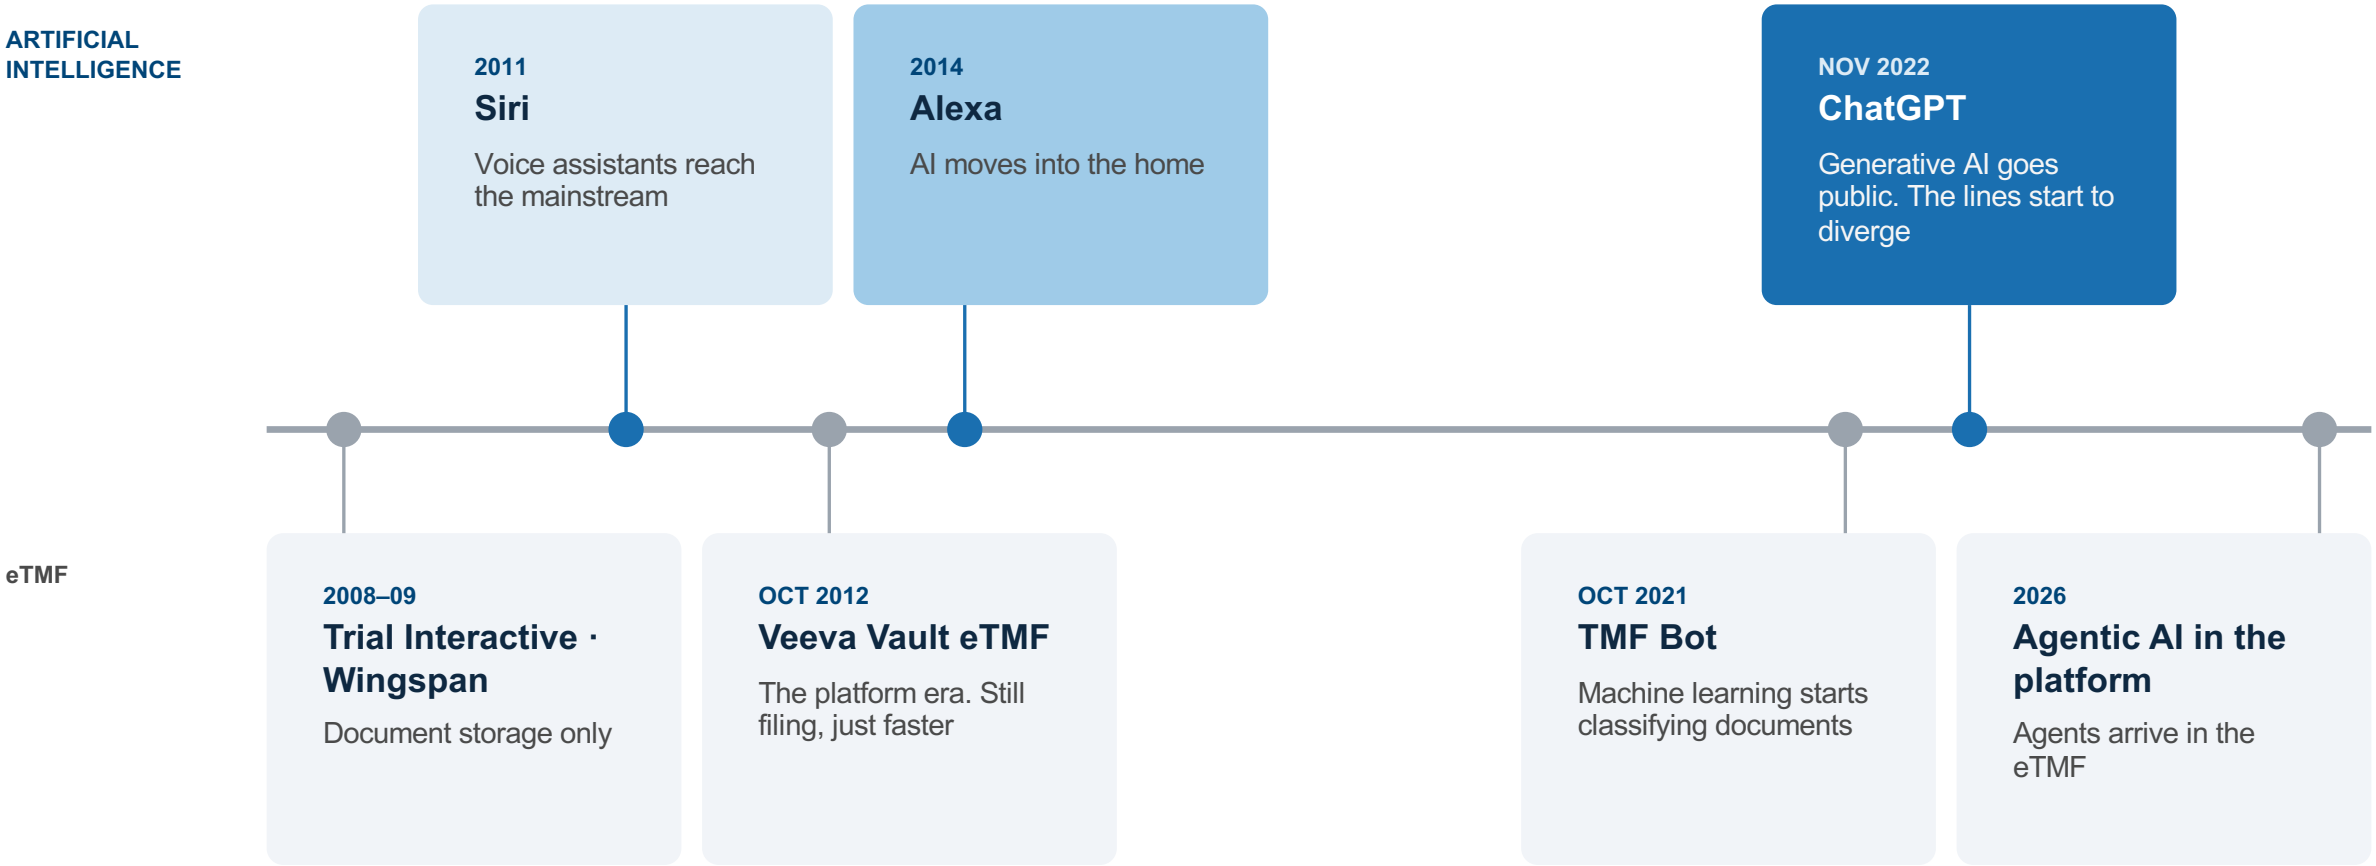
04 The way forward

The awkward question: should we teach this?
The practical takeaway: building the plan

cdisc®

ARITHMOS

We were running in parallel for years



Since 2022 the AI line has shot ahead while the eTMF line has stayed roughly flat.

cdisc®

ARITHMOS

Why this matters to us specifically

Huge datasets

Our work is large-scale document and data review, which is exactly what these tools are good at.

Knowledge bases

Years of SOPs, guidance and precedent sitting in systems nobody can search properly.

The regulatory angle

Everything we produce has to be attributable and inspectable. AI has to fit that, not the other way round.

FIRST OF ITS KIND

Purolea Cosmetics Lab, 2 April 2026

The first cGMP warning letter to cite the misuse of AI.

- AI agents generated drug product specifications, procedures, and master production and control records
- No process validation: the firm said it didn't know it was required, because the AI agent never told it
- The finding isn't that AI was used. It's that nobody properly reviewed what it produced

A GMP citation, not GCP, but the principle transfers directly.

Learning from your own mistakes is fine.
Learning from someone else's is better.

Five levels, one team

1	**Manual**	Calculator-in-Excel. Hasn't really engaged with AI; everything is still done by hand.
2	**Dabbler**	Has tried ChatGPT a handful of times for something low-stakes. Wouldn't trust it with real work.
3	**Practitioner**	Uses AI daily for real tasks like summarising and drafting. Hasn't built or configured anything.
4	**Builder**	Goes beyond prompting: custom skills, basic automations, tools stitched together. Not yet agents or MCPs.
5	**Architect**	Builds their own agents, connected into systems via MCP, with interactive dashboards on live data.

Beyond Level 5: AI research and development. Real, and deliberately outside this talk.

What happens when the levels collide

1

The Veeva agent

- Built to invoke actions, like checking a document or a value, not to hold a conversation
- Levels 1–2 won't touch it: still AI, still unfamiliar
- Levels 4–5 may find it limiting: no follow-up questions, no rabbit holes
- Competency alone doesn't predict comfort. The tool has to match how people expect to use AI

2

The sign-off gap

- A Level 4 builds a TMF health check on live data: heat maps, radar charts, real insight
- The Level 2 who approves it can't meaningfully challenge the method
- It gets signed because it looks impressive and comes from a trusted colleague
- The GMP finding one layer removed: "my colleague's dashboard said it's fine"

3

Excel or a bespoke dashboard

- Same job: deep-diving TMF health because system reporting isn't enough
- Excel pivots plus experienced commentary: explainable in thirty seconds
- An interactive dashboard with drill-downs and anomaly flags: more powerful, possibly a black box
- Both valid, both compliant, very different transparency

Before anyone says "just follow the process": process only ever covers the backbone.
Everyone brings their own style within it. AI competency is the newest, most visible version of that.

So what does training actually look like?

One programme for everyone loses the room at both ends, and a rigid track per level loses people for the same reason. Three layers instead, built from the foundation up.

2

TOOL-SPECIFIC · WHATEVER YOU'VE DEPLOYED

What changes tomorrow morning

Copilot, Claude Code, Veeva's agents: what the tools you actually use can and can't do in your context. Grounded and practical, not AI theory.

3

OPTIONAL SKILLS TRACK · ELECTIVE

Go deeper if you want to

Prompting technique, how tokens and context work, what the models are doing under the hood. Open to anyone ready to move up the scale; forced on no one.

1

MANDATORY FOUNDATION · EVERYONE, EVERY LEVEL

Compliance, risk and governance

Non-negotiable, because it protects the team and the organisation, not just the individual. Starting skills: check the output · know what data goes where · spot a tool out of its depth · know when review or validation is required.

Layer 2 ties back to the Veeva agent: much of that friction was people not knowing what the tool was built to do.

Role matters as much as level

Document-level user

AI recommends where to file something. That is the whole interaction.

TMF Manager

New ways of reviewing and dissecting data: the pillar-three problem.

System Owner

Accountable for a system that has started making judgement calls.

System owners carry the validated state. Embedded AI, running in the background and cross-referencing, changes what that job is.

This is structural, not just a training gap

Your system owner knows the regulations and the TMF inside out, and sits at Level 1 or 2 on the AI scale.

- Should the role split into validated-state ownership and AI governance ownership?
- Can one person genuinely hold both going forward?
- Is Level 1 or 2 still a viable place for a system owner to sit?
- “Can you explain what this feature is actually doing?” is now an inspection question

Resourcing reality: in many organisations the TMF function is a team of one.
That person has to be the regulatory expert and the AI-literate owner. Is training alone enough to close a gap this size?

We have already done this once

ALREADY HAPPENING

Not a prediction

Job descriptions in this industry are asking for AI competency now. This is a current reality, not a forecast.

WE'VE BEEN HERE BEFORE

Paper to electronic

People knew every binder in the filing room. That mastery didn't protect them when the TMF went electronic. The differentiator became who adapted.

AI isn't your competition.

Your competition is who it always was: colleagues who adapted.

That shift was learnable, and so is this one.
Progression will favour people comfortable with both the system and the AI built into it.

Should we even be the ones teaching this?

This moves fast. New models arrive every few months, and the same prompt can behave differently after one lands.

THE BASICS

Anyone can teach these

- What a model, a token and a prompt actually are
- What a good prompt looks like, and why the same prompt can give a different answer
- L&D, vendors and online courses do this well, and keep it current

HOW IT'S USED IN OUR WORK

This part is ours

- When a new model lands, most people adjust the prompt without thinking twice
- In a regulated process, changing the prompt is changing the system configuration
- Only people who know GxP will treat it that way

Who teaches it? An AI-curious CTA who knows the job beats an AI expert who hears "GCP" and thinks Google Cloud Platform.

We're not teaching AI.

We're teaching AI in the context of the TMF, and that second half is the part nobody else can do.

Building the plan

1

Map the spread

Anonymous self-assessment, 1 to 5, sense-checked against real output. Expect a spread, not a cluster.

2

Overlay the roles

Plot maturity against role. Where accountability outruns competency, you've found a risk, not a training need.

3

Build the layers

Mandatory foundation for everyone, tool-specific training for what you've deployed, and an optional skills track.

4

Re-baseline

Reassess when the tooling changes, not annually. It will change long before the calendar does.

And TMF training changes too: less about where a document goes, more about why.
So people can tell when the AI has got it wrong. Checking AI output becomes a core TMF skill.

The goal isn't to get everyone to Level 5.
It's to make sure nobody is accountable for something they can't explain.

Questions?

Same Team, Different Worlds

Rob Jones · Arithmos Tech
rob.jones@arithmostech.com
www.arithmostech.com

