capSpire



TRADING AI

10 REASONS WHY TRADING AI PILOTS STALL

Ten things that go wrong before anyone trains a model. None of them are the model.

Most firms engaged in commodity trading activities have now funded at least one AI pilot. Most are still pilots. The reason is rarely the model, the vendor, or the team. It is that the data underneath was never built to carry it, and the work of fixing that sits between the ETRM team, IT and trading operations, so nobody owns it.

RECOGNIZE THREE OR MORE? THE PILOT ISN'T THE THING TO FIX.

WHERE THE MONEY GOES

01	**The pilot that demos well and goes nowhere**	The use case gets picked because it looks good in a demo, then stalls because the data, access and ownership underneath it isn't ready.	→ Choose a first use case that forces you to build something you will use again.
02	**Expecting the reporting warehouse to cope**	It answers the questions you already knew to ask. AI asks the ones you didn't.	→ Keep the detail. A single summarized layer throws the answers away.
03	**Automating on top of unstable data**	Automation built on unstable data is not automation. It is a maintenance burden you carry indefinitely.	→ Stabilize the data services first. Automate second.

WHERE THE RISK SITS

04	**Building without people who know both sides**	Data engineers who don't know the trade lifecycle misread amendments, books and curve versions. ETRM experts who don't build data services struggle to make the answer repeatable. AI that works needs both on the same team from day one.	→ Put ETRM experts and data and AI engineers on the same team from day one.
05	**Forgetting that trades change**	Trades amend and positions restate. A model trained on one snapshot is confidently wrong.	→ Being able to ask what the position looked like on any given day, not just today.
06	**Leaving permissions until after the pilot**	Copy data out without book and counterparty permissions attached and you have a compliance exposure.	→ Attach permissions to the data so they travel wherever it goes.

WHY IT NEVER GETS FIXED

07	**An extract is not a data model**	CTRM/ETRM tables copied into a lake are the same tables, just further from the people who understand them.	→ Build one modeled layer where a trade means the same thing to every consumer.
08	**Building the same connections, one at a time**	Everyone agrees this was a mistake last time. It gets repeated anyway, one interface at a time.	→ New systems connect to the foundation, not to the ETRM.
09	**Underestimating the reference data**	Counterparties, instruments, locations, units. Unglamorous, never funded, and the reason nothing reconciles.	→ Own and resolve it once, upstream, before anyone builds on top.
10	**Nobody owns the layer in the middle**	The data foundation sits between the ETRM team, IT and trading operations, so none of them fund it.	→ Name one owner. Fund one budget line.

WHY CAPSPIRE

Trading technology consulting specifically for the energy and commodities sectors has been our work since 2009, across every major C/ETRM platform. More recently we have been helping clients with these exact challenges with their AI plans. We sell no software licenses, so our recommendation follows what your operation needs.

30 MINUTES WITH PEOPLE WHO HAVE ALREADY FIXED THIS. NO PITCH.

capSpire.com