

**We don't just move data;
we move human possibilities.**

And in the agentic era, we move consequences.

Why trust, accountability and governance are prerequisites for letting AI act on your behalf —
not afterthoughts you fix once something has already gone wrong.

HARVEER SINGH

Author, *When Data Moves* · Former Chief Data Officer

A Learning Journey

CHICAGO · 2026

PICKING UP WHERE WE LEFT OFF

You've just heard that business-ready data
ready data doesn't happen by accident.

Here's what happens when it doesn't — and something acts on it anyway.

accident.

The one thing to take with you

*An AI that looked at bad data, and an AI that acted on it, are two different problems
— and governance is the thing that separates them.*

An agent can be responsible for what it does.

Only a human is accountable for what it means.

DECLINED

Zoom out.

It's a hand holding a card.

Zoom out.

She's standing at a gas pump.

Zoom out.

There's a child asleep in the back seat.

Zoom out.

**The station is closed.
The lights are off.**

Zoom out.

It's a highway. It's dark.
There is nobody for miles.

FROM HER, TO YOU

Now make her your mother.
mother.

Your sister. Your daughter.

Somewhere, a system said one word about her.

And that was the whole story it told.

I built her out of a hundred real ones.

A version of this happens somewhere every night.

HERE IS WHAT THE SYSTEM ALREADY KNEW

None of this was missing.

LOCATION

01

A highway station, far from any address she's ever used.

TIME

02

After midnight — outside every every pattern she has.

HISTORY

03

Twelve years. Never a late payment. Never a dispute.

CONTEXT

04

A small fuel purchase. Attempted Attempted three times. Then a a fourth.

*The data that could have told the whole story was already in the building.
We just chose to look at one field.*

DECLINED.

AND HERE IS WHAT I AM NOT SAYING

I am not asking you to *approve the transaction.*

That is not the lesson. Sometimes the decline is correct.

Fraud is real, and I have spent a career fighting it.

I'm asking why nobody sent help.

A text with the nearest open station

A number to call

An offer to cover twenty dollars of fuel

Not the next best action.

The next best human action.

Stop asking "what can we sell next?" and start asking "what does this person need right now?"

I build for humans.

So I start with why.

Every system I have built moved somebody's medicine, tuition, payroll, or rent.
The pipes don't judge — but the people who build them decide what flows through.

*"I've run this exact fight from inside a bank
with 40 million customers."*

There are five different people in this room.

01

If you run the business

*What it costs when nobody checks
—
and the question to ask before you
approve the next agent.*

02

If you architect systems

*Where trust and accountability live
accountability live as controls, not
not principles.*

03

If you build the pipelines

*Why the data contract is the
safety system.*

04

If you build on SAP or Reltio

*What has to be true underneath
underneath before anything acts
acts on its own.*

05

If you're here for the ideas

*We are not moving records.
We are moving what happens to
people.*

One promise to all five: you'll leave with something you can act on Monday.

SO WHAT WOULD HAVE MADE THE DIFFERENCE

That system had the data.
It just couldn't trust it.

*Trust is not a feeling about your data. It is not a dashboard that says 98% complete.
Trust is the engineered property that decides whether a system is allowed to act on something —
without a human standing behind it.*

This is the prerequisite.

Everything else in the agentic era sits on top of it.

Think of it as a farm-to-table restaurant.

The menu only works if the produce does.

FARM

Sends spoiled produce.
— a quiet, structural failure.

CHEF

Is world-class. And useless.
— sharper knife, bigger kitchen, Michelin. None of it.

DINER

Blaming the restaurant.
— never the soil. Never.

*AI is the chef. Your data is the farm.
The diner never blames the soil.*

Five things that have to be true before an agent acts.

agent acts.

01

CONTEXT

The entity, not a row

*It understands the entity it's acting on —
a woman with twelve clean years, not a row in a table.*

02

LINEAGE

**Every decision reconstructable
reconstructable**

Every decision is reconstructable: what data, what model, what rule. Someone can follow it backwards.

03

GUARDRAILS

Enforced before day one

*The bounds are encoded before day one.
Enforced rules, not slide-deck principles.*

04

A DEFINED 'WRONG'

**You named failure before
deployment**

You named failure before you deployed. You cannot catch what you never defined.

05

AN ACCOUNTABLE OWNER

**A named human answers for
the call**

*A named human who answers for the call.
Autonomy without ownership isn't innovation — it's negligence.*

Miss any one of them and you don't have trusted data. You have data you're hoping about.

We used to fight about databases.

databases.
Now we fight about models —
same religion, wearing a lab coat.

Both are faith arguments. Neither one produces a number.

And once a system is allowed to act,

you cannot govern it on faith — only on evidence.

A background image of Tom and Jerry in a kitchen. Tom is on the left, looking angry with his mouth open. Jerry is on the right, looking happy and holding a piece of cheese. The kitchen has wooden cabinets and a table.

Tom never catches Jerry.

He was never supposed to.

If the cat caught the mouse, the show would end —
so the format depends on the problem never being solved.

Enterprise data has worked the same way.

Every few years we buy a faster cat.

Warehouse	---	Lake	---	Lakehouse	---	Agent
-----------	-----	------	-----	-----------	-----	-------

Why is the mouse in the kitchen?

BUT HOW DO YOU MEASURE ANY OF THAT

*You cannot govern what you
cannot price.*

Every one of those five elements sounds obvious in a room like this.

None of them survive a budget conversation
until somebody can say what it costs when they are missing.

**So three weeks ago, I sat down to
build it.**

*And then I got it spectacularly wrong —
which is the most useful thing I can give you
today*


The Amplified Cost of Errors

THE QUESTION

Does adding AI + Data + Humans
reduce risk .. or multiply it by 10x?

AI ERROR

How far are we from expected outcome?



\$17,500,000,000

of exposure. From one bad credit record.

Five agents. A hundred thousand inferences a day.

One bad record, multiplied by everything that ever touched it.

That is not a big number.

That is a broken formula wearing a big number as a costume.

My formula could not tell the difference between
two completely different things.

CASE A · INPUT

*The AI looked
at bad data.*

- reads a record.
- flags a confidence score.
- stops here.

CASE B · AGENCY

*The AI did something
because of it.*

- turns a record into a decision.
- an action.
- a consequence.
- with nobody in between.

*An agentic system doesn't just process a record.
It turns a data error into a decision, and a decision into a consequence —
in one motion, with nobody in between.*

Not a value statement. A missing variable.

\$17.5B → **\$108**

Two new terms in the formula

$p(\text{action})$ the probability an output actually leads to a real action —
and the deduplicated count of distinct decisions it altered, not how many times a model looked at it.

*The difference between a number that sounds impressive
and a number I can defend to a CFO, line by line.*

AND THEN I CHECKED MY OWN CORRECTION

I had already caught the big one.

I still got the small one wrong.

HEADLINE vs OWN FORMULA

~~\$17.5B~~ vs **\$17.15B**

*Run my own formula as written and it is seventeen point one five —
a three hundred and fifty million dollar arithmetic slip,
inside the very example I was using to warn you about inflated numbers.*

And this is the part that should worry you.

THAT WAS MY VERSION · HERE IS YOURS

DETAILS CHANGED — THE FAILURE IS REAL

Somebody chose a data type.

A payment platform stored transaction amounts as text.

*It passed every test, shipped, and ran correctly for eighteen months —
because text sorts alphabetically, and nothing had enough pressure on it to notice.*

“\$100” sorts before “\$2”

(MONO TYPE)

2:47 PM
SORTING FAILURE

3:15 PM
COLLAPSE

3.2 HRS
DOWN, WORST DAY OF THE YEAR

No contract. No lineage. No defined 'wrong'. Trust was assumed — never engineered.

Same lesson, eighteen months instead of three weeks.

Finding Nemo is a story about data.

Nobody notices, because it has fish in it.

01

DORY

Brilliant in the moment, no memory.

Reads the whale script.

Forgets the plan.

Most AI in production today: capable per query, amnesiac across them.

02

THE TURTLES

Intelligent, and they remember.

Generational context.

The whole current, mapped and passed on.

This is what you are trying to build.

03

MARLIN

The one with everything at stake.

Neither of them finds Nemo.

The parent swims the distance himself.

They got him there. They did not go for him.

Let that land.

You can buy Dory. You can build the turtles. Somebody still has to be Marlin.

My name does not matter.

My actions do.

*And now my agent's actions do too —
because there is nobody left to
blame.*

Governance isn't a gate you pass through once.

*It's a habit —
especially right after you've congratulated yourself
for catching the last mistake.*

A human is always **Accountable**.
An agent is **Responsible** for execution.

Accountability never delegates.

Five things make data trustworthy. Three controls are where you enforce them.

01

The contract

**Context & lineage,
enforced at the boundary.**

*So "\$100" can
never be text again.*

02

The gate

**Guardrails & your defined
'wrong'
declared before deployment.**

*Wired to stop it.
Not narrated after.*

03

The name

**One accountable human,
in the system of record.**

*A name that answers for
what the agent did.*

*And one more, for the woman at the pump:
give your systems more than one thing they are allowed to do.*

If you build on SAP or Reltio : the contract is your data model and survivorship rules / the gate is your match thresholds and workflow approvals / the name belongs to a domain owner, not to the platform.

Many agents, one dataset. Now you have have arms.

01

Same data, many angles

Fourteen agents read one dataset at once — risk, fraud, servicing, pricing. Each works a different dimension of the same question.

02

Composed, not averaged

*Their findings assemble into cause cause and effect no single graph could produce.
This was not possible three years ago.*

03

*An argument, not an answer
What comes back carries its evidence evidence — in the time it used to take take to get the meeting in the diary. diary.*

04

Cheap to ask again

*Re-running the question costs minutes, not a quarter.
You can afford to ask before you commit instead of explaining after.*

05

You still act

The judgement stays where the where the accountability already sits. Nothing has been taken away from you.

Analytics gave you eyes. This gives you reach: fourteen directions on directions on one problem before you decide, instead of one direction after the fact.

Go test your own foundation.

I got tired of watching the same failure in a different company every year —
a woman declined on a dark highway, money stored as text,
a formula that says seventeen billion dollars. So I built the check I wish I'd had.
DataPulse is free and open source: fourteen dimensions across your data foundation and AI readiness.

01

This week

*Run DataPulse against one domain.
Not the whole estate — one.*

02

Then ask

*Which of the five is missing?
Context, lineage, guardrails, defined 'wrong', owner.*

03

Then decide

*Whether anything in that domain has earned
earned the right to act on its own yet.*

DATAPULSE · FREE & OPEN SOURCE

github.com/harveer10x/DataPulse

Somewhere in your architecture is a field nobody ever defined.

Nobody validated it. Nobody owns it.

*It's passing every test today — and it's the one your agent will act on,
while somebody's mother stands at a pump in the dark.*

WHAT'S YOUR VARCHAR?

WHERE THIS GOES NEXT

Trust it. Own it.

*That's when context can finally turn it
into enterprise intelligence.*

WHEN DATA MOVES

Building technology that remembers we're human.

HARVEER SINGH

Thank you.

THE BOOK

Signing tonight · 4:30 – 6:00 PM

Extract wisdom. Transform lives. Load the future with hope.

harveer@rizzwireless.com · [linkedin.com/in/harveer](https://www.linkedin.com/in/harveer)