

From Rearview to Radar:
AI-Powered Risk Sensing

Mike Roberts

Vice President, Analytics | Riskonnect, Inc.

2026 Konnect

JAN	**A control test fails in a shared service centre.** Logged, remediated and closed. Nobody did anything wrong.
FEB	**An issue closed in November is re-opened.** Different owner this time, and no link back to the first closure.
FEB	**A supplier assessment score slips from 3 to 4.** Nothing breached a threshold, so nothing was notified.
MAR	**A near miss is logged in a different system.** Different team, different taxonomy, same site.
MAR	**Three business units rate the same process higher.** One point each, across three separate assessment cycles.

APRIL: the heat map turns red.

Five signals, four different systems, and nobody whose job it was to look at them together.

5 of 5

signals captured on time and in the right place

0 of 5

reviewed in context of the others

This is not a data collection problem.

It's a connection problem.

Somebody did own the cross-domain view, but the signals sat below the threshold that would have alerted them.

Designed for Assurance, now asking for Detection.

ASSURANCE

The job the process was built for

Assess most risks annually, high-profile risks quarterly

Designed above all to be defensible

Optimised for accuracy after the fact

Tolerates being slow

Answers: can we prove this was managed?

DETECTION

The job now being asked

Continuous and specifically between cycles

Designed above all to be timely

Optimised for early warning

Tolerates being wrong

Answers: what changed since we last looked?

Risk teams are not backward-looking by choice.

These are different jobs with different tolerances, but almost every process in the function was designed for the one on the left.

The annual BIA survey is an assurance instrument being asked to do detection work.

2026 KOnnect

Boards have changed. The plumbing hasn't.

Over half

of risk and insurance professionals lack meaningful system integration. Their technology sits across disconnected systems and spreadsheet-dependent processes.

Airmic Risk Technology Survey 2026, "Mind the Gap"

84% vs 10%

of directors have significantly changed their approach to escalating risk. Only one in ten is using AI tools to manage the complexity that creates.

2026 What Directors Think, Diligent & Corporate Board Member

3 in 4

of technology users report better reporting as the demonstrable benefit. What risk technology has most improved so far is the rearview.

Airmic Risk Technology Survey 2026, "Mind the Gap"

Without integration, everything downstream is limited: Data quality, Analytics and AI.

2026 KOnnect



Implementing a Radar… Use the signals you already have

2026 KOnnect

Four kinds of signal, already in your data

1

FREQUENCY

How often something recurs.

Repeat incidents logged under different labels. Issues re-opened after closure. Control tests failing at the same point in every cycle.

2

DRIFT

Direction, not position.

One point of movement is invisible on a heat map. The same movement in three business units in the same quarter is not.

3

LATENCY

How long things take.

Time-to-close lengthening. Attestations becoming overdue. Control environments get slower before they get worse.

4

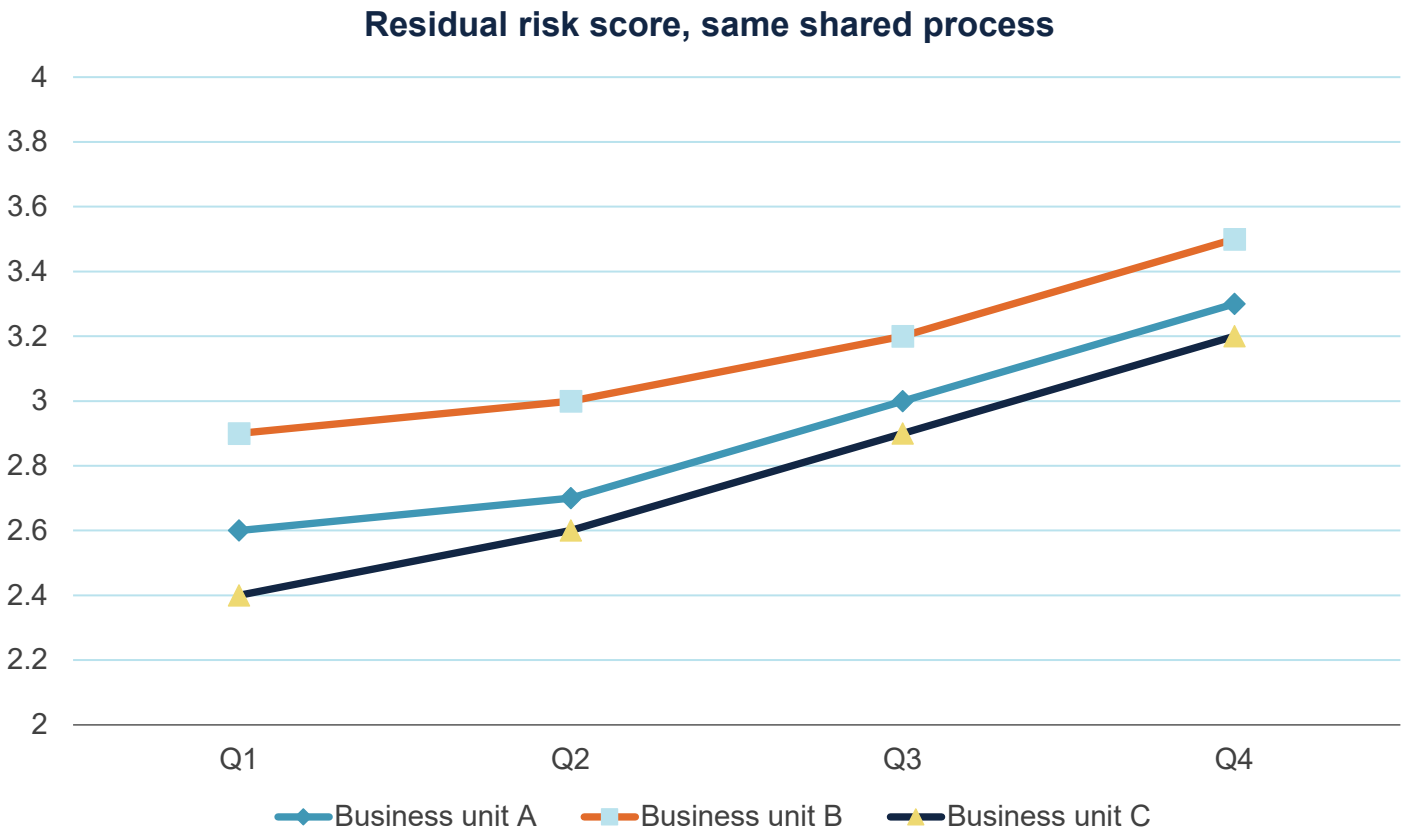
SILENCE

Expected data that stops arriving.

A business unit that stops logging near misses hasn't become safer. Absence is the signal a human reviewer will never spot.

None of this requires any new data collection, and none of it requires predicting anything.

Drift: three cycles, three business units



Illustrative figures, for shape rather than precision.

What each owner sees

A single risk moving 0.7 over a year. Within tolerance, no threshold exceeded, nothing to escalate.

What nobody spotted

The same process, rated by three independent teams, moving in the same direction at the same rate across three separate assessment cycles.

Correlated drift is a far stronger signal than any of its parts. And you can compute it, which means nobody has to spot it.

Look for the same pattern elsewhere, e.g.,

Criticality ratings for comparable processes, compared across every BIA in the programme.

2026 KOnnect

Start here: the query you could run on Monday

THREE FIELDS YOU ALREADY HAVE

- Date the action or issue was raised
- Date it was closed
- The owning function or business unit

ONE CHART

Median days to close, by function, by quarter. Four quarters is enough to see a shape.

WHAT A BAD SHAPE LOOKS LIKE

- A function whose median climbs three quarters running
- A function whose median is flat but whose spread is widening
- A function where the open count is stable but the items are ageing

None of these proves anything, but it's a conversation you would not otherwise have had.

No AI. No new fields. No project. If the line bends, you have a question worth asking before anything fails.

Where the reports run out…

IT ONLY FINDS WHAT YOU SUSPECTED

Every threshold rule encodes a hypothesis somebody designed. Sometimes the crucial signal is the one nobody had considered

IT NEEDS THE DATA TO BE CLEAN

Grouping by function assumes function is populated the same way everywhere. Filtering on a category name misses the three teams who chose a different one.

ZERO ROWS LOOKS LIKE GOOD NEWS

A query with the wrong category name returns nothing. Nothing is indistinguishable from nothing to worry about, and no one investigates an empty result.

THE REAL QUESTION ISN'T A QUERY

"Should I be worried about the shared service centre?" isn't one question. It's dozens, across three domains, and a human analyst's first response would be "worried about what?" A filter has no way to ask that.

A report answers the question you thought to ask. A Risk Radar asks you one.

From Signal to Radar to Recommendation

Operational signals

Related records, contextual triggers, and thresholds crossed across the register.

Materiality test

Configured by you, not inferred by the model. You can define what counts as a material change.

Reassessment recommendation

Routed to the risk owner, with the evidence that triggered it attached.

The exposure gap

Scheduled assessments still happen. Signals are what happens between them, when you're running on stale intelligence.

The Risk Signal Agent does this today…

It does not re-score your risk, and it does not predict an event. It notices that something material changed and tells the person who owns the decision. The judgement stays where it was, with you.

2026 KOnnect



Connected risk is the multiplier

2026 Konnect

Connection runs in three directions

1

ACROSS DOMAINS

GRC, Resilience and Insurable Risk often describe one event from three perspectives.

2

ACROSS THE BOUNDARY

A supplier's finances deteriorate, or curated risk knowledge from outside your register moves, and a risk you own needs reassessing.

3

ACROSS TIME

Some patterns exist only in the history. Incidents recur under different labels; criticality ratings drift across successive BIAs.

The AI's job: decide that three differently-named risks are one risk and read the connections that live in the descriptions rather than the schema.

2026 KOnnect

Across domains: one provider, three registers

GRC

A provider's assessment score slips a band at annual reassessment. An audit exception is noted and accepted with compensating controls. The quarterly review of their privileged accounts is running late.

All three handled correctly.

RESILIENCE

Provider is an upstream dependency for two processes rated critical in the last BIA. The recovery objective assumes their availability, which has never been tested jointly.

The document is accurate and static.

INSURABLE RISK

Contingent business interruption sits at a sub-limit well below the main limit, behind a waiting period of eight hours. A short outage at their end recovers nothing.

Nothing moved. It was always this.

One cause: the provider changed its subcontractor.

Three teams, all doing their job correctly. Three registers, none of them wrong, but nothing records that change.

Any one of those signals is noise. Together they describe a supplier you now depend on differently from the one you assessed.

Why risk data is difficult territory for AI

SPARSE	The events you most care about are rare by design. You're trying to learn a pattern from a handful of examples.
SUBJECTIVE	A risk score is a judgement rendered as a number. Two competent people score the same risk differently.
SELF-REPORTED	A lot of risk data is assessed by people with a legitimate interest in how the answer looks.
INCONSISTENT	The same risk is named three different ways in three business units, and all three are defensible.
COUNTERFACTUAL	Much of the data describes things that didn't happen. There's rarely a ground truth to train against.

Four foundations

1. **Use one name for one thing**
   Taxonomy consistency across business units. Without it, correlation is guesswork. You can't count recurrences of something with four different names.

2. **Keep the history**
   Drift and latency are invisible unless state changes are timestamped and retained. Most registers store where a risk is and throw away how it got there.

3. **Own the threshold**
   Materiality is a business judgement. Somebody has to make it, and it isn't the model. Expect to change it twice in the first year.

4. **Close the loop**
   Keep the trail of what was recommended and what was actually done. That record tells you which alerts mattered, and it's where cross-incident patterns come from.

Treat the AI as a control

WHAT IT IS

Prompts, model versions, materiality thresholds, the scope of records read. Configuration, and it governs behaviour more than the model does.

WHAT IT NEEDS

A named owner. A change log. Approval before modification. Gated access. Evidence that it was tested.

WHY IT MATTERS

New models are always becoming available. A prompt that behaved one way in March might behave differently in July.

You don't need a new discipline for this.

If an auditor asks why a reassessment fired on a given date, you need the prompt and the model version as they were that day. That is an evidencing problem, and you already own the machinery for evidencing problems. Add the AI configuration to the control inventory you maintain.

The catch: if it is a control, it has to be tested for effectiveness.
Testing an LLM's output is difficult, but not impossible.

Not a crystal ball, just surfacing what you didn’t see

REPORT
What happened

Defensible, necessary, and structurally too late to change the outcome.

SENSE
What changed

Correlated signals, a materiality threshold, a recommendation between cycles. Available now.

ANTICIPATE
What is likely

Simulating disruption across the full dependency chain, upstream and downstream, instead of forecasting the event itself.

Prediction in risk is difficult today mostly because the input data won't support it.

Simulation takes a different route. If you know what depends on what, you can model a disruption and see its blast radius upstream and downstream without forecasting the event at all. Sparse data limits prediction. Mapping dependencies is a different problem.

Every one of the four foundations is groundwork a forward-looking capability needs anyway. Build it for sensing and you've built it for whatever comes next.

What does your organization already know that nobody has looked at together?

Not what data you're missing. What you already have.

QUESTIONS

2026 KOnnect

CONNECT WITH ME

Mike Roberts

mike.roberts@riskonnect.com

riskonnect.com

in/mike-roberts-59b0b42/

2026 KOnnect


THANK YOU
2026 Konnect