

GLOBAL CONFERENCE 2026

FROM CHAOS TO COMMAND:
RISK LEADERSHIP IN INDUSTRY 6.0

**DATA TO DECISIONS:
ENGINEERING DECISION
INTELLIGENCE IN THE AGENTIC ERA**

Kathiravan Manoharan
Global VP of Analytics
Collinson

IERP® GLOBAL CONFERENCE

Data to Decisions: Engineering Decision Intelligence in the Agentic Era

Kathiravan Manoharan · Global VP of Analytics · Priority Pass | Collinson

19th AUGUST 2026, KUALA LUMPUR, MALAYSIA



**Every organisation in this room already has data,
dashboards, and increasingly, AI.**

That was never the hard part.

THE THESIS

Insight is now abundant. Decision quality is still scarce.

THE FRAMEWORK

Five structural shifts define the intelligent enterprise.



HOW WINNING ORGANISATIONS WILL COMPETE

**Leaders of the future will execute TRUSTED
decisions at SCALE.**

THE REAL DIFFERENTIATOR

Everyone has the inputs now. Almost nobody has engineered the output.

WHAT EVERY ORGANISATION HAS

- Dashboards and reports
- Predictive models
- Increasingly, AI agents
- A growing sense that none of it is quite enough

WHAT STILL SEPARATES THEM

- A named decision the model is meant to improve
- A trusted signal — not just an accurate one
- A system engineered for judgement, not just output
- Someone accountable for the decision, not just the model

DEFINITION

Decision Intelligence, precisely defined.

The discipline of engineering the decision itself. It's about mapping who decides, what evidence is required, and how AI and human judgement combine at each step. So choices get better and faster on purpose, not by accident.

Not just better BI

BI tells you what happened. It can automate insights & signals. Decision Intelligence engineers what happens next.

Not just automation

Automation executes a known path. This engineers the judgement behind the path itself.


THE AUTOMATION LINE

Not every decision deserves the same treatment.

Some should be automated in milliseconds. Others should never leave a human's hands.

WHERE TO DRAW THE LINE

Automate low value, low risk. Protect high value, high risk.



WHO DRAWS THE LINE

Whose risk, whose value — and who actually owns that call.

ENTERPRISE'S RISK & VALUE

Loss exposure, regulatory penalty, reputational damage, cost to serve.

CUSTOMER'S RISK & VALUE

Access denied, time lost, dignity, financial hardship — and they often point in opposite directions from the enterprise.

THE DECISION OWNER

A named accountable role per decision domain — not a standing committee. Risk holds a hard veto on the risk classification. Finance or Commercial owns the value threshold. The Decision Owner is who a regulator, a board, or a court asks first.

TRUSTED SIGNALS

Trust is not a feeling. It's a test.

A track record tells you it worked before. It doesn't tell you it's still working.

THE TRUST TEST

Three tests a signal must pass before it earns trust.

1

Calibration

Stated confidence matches actual accuracy, over many decisions — measured with a calibration score, not a gut feeling.

2

Explainability

Any single decision can be audited after the fact, in language the accountable human already reasons in — not just the model's global logic.

3

Stability

The signal holds up as conditions drift. The world the model was trained on keeps moving; the test has to keep running.

WHERE AGENTS CHANGE THE EQUATION

Not every use of AI is agentic. Here's the difference that matters.

THREE TIERS, NOT ONE

A dashboard informs. A model scores. An agent acts.

1

Dashboard

Human reads, human decides, human acts. Hours to days between signal and action.

2

Model / Rules Engine

System scores or flags a pre-defined path; a human executes it. Fast, but brittle — it only handles branches someone explicitly coded.


3

Agent

System reasons across steps and executes within policy, escalating only on exception. Generalises to situations nobody pre-coded, in seconds.


BRINGING IT TOGETHER

One loop: signals, decided, acted, and fed back.



EXAMPLE 01

Fraud detection, engineered end-to-end.



BEFORE: HOURS TO DAYS


An analyst manually works the queue, cross-referencing device, account, and case history before deciding.

NOW: UNDER 60 SECONDS

For transactions in the automate-fully quadrant — low value, low risk. Anything ambiguous still escalates to a named fraud analyst.

RISK TOLERANCE AND THRESHOLDS

Automation maturity is a curve, not a switch.



This is the same automation line from earlier — moving over time, not set once.

MEASURING DECISION QUALITY

Don't just measure the outcome. Measure the decision.

Outcome-only scoring rewards lucky bad decisions and punishes unlucky good ones.

TWO AXES, NOT ONE

A good decision and a good outcome are not the same thing.



Calibration at scale, review by exception — not by volume.

1

Calibration Score

Stated confidence versus actual hit-rate, computed automatically on every decision — a Brier-style score, not a person checking each one.

2

Stratified Sampling

Heavier human review where value and risk are both high. Thin, low-frequency sampling where both are low — same logic as manufacturing QA.

3

Drift-Triggered Review

When live calibration diverges from baseline, that divergence itself triggers a review — nobody has to notice or remember to look.

Decision quality becomes the metric that matters most.

Decision Quality

The single measure that captures whether the entire transformation is working.

Speed

Time from question to validated decision.

Accuracy

How often the decision proves right in hindsight.

Trust

How readily leaders act on it without re-checking.

ACCOUNTABILITY AT MACHINE SPEED

An agent can be responsible for the action. It cannot be accountable for the outcome.

THE LINE THAT DOESN'T MOVE

The agent owns the action. A named human owns the outcome.

THE AGENT

- Executes within reviewed policy bounds
- Logs the full rationale for every action
- Escalates automatically on exception

THE ACCOUNTABLE HUMAN

- Approved the policy before go-live
- Owns the sampling and audit methodology
- Answerable for it to a regulator, board, or court

BNM'S OWN GOVERNOR, JULY 2026

As AI adoption accelerates across Malaysia's financial sector, banks must be able to explain outcomes and remain accountable for how AI is used — boards and senior management should treat it as a core agenda item. He stressed that accountability for AI-driven decisions must remain with human leaders.

The decision log has to block, not just record.

1

100% Capture

Every decision logged automatically — never sampled. Logging is cheap, automated data capture; only review is sampled.

2

Blocking Approval

Decisions above a defined value or risk threshold cannot proceed without a completed log entry — no exceptions, not optional.

3

Feeds the Calibration Loop

The log isn't filed away. It's the raw material the sampling, review, and calibration system runs on every day.

DECISION INTELLIGENCE BELONGS IN THE BOARDROOM

The next competitive advantage is not better data. It is better decisions.

What Does this mean for THE BOARD and THE CEOs?

1

Capital Discipline

Stop measuring AI spend by adoption. Measure it by decision quality — the same rigour you already apply to any other capital allocation.

2

Workforce Transition

Some roles will not survive this shift. Naming that honestly, early, and with a redeployment plan is what keeps trust intact through it.

3

Accountability at Machine Speed

The moment a decision is made by an agent, not a person, someone in this room is still accountable for it. Design that in before it's tested — not after.

CLOSING

In the age of Agentic AI, every organisation will have data. They will be distinguished by the quality of the decisions they make.

GLOBAL CONFERENCE 2026

FROM CHAOS TO COMMAND:
RISK LEADERSHIP IN INDUSTRY 6.0

THANK YOU



SCAN THE QR CODE
TO KNOW MORE ABOUT US!

enquiry@insterp.com
www.insterp.com