

BOARD BRIEFING

The Ten Discussions Where KAIIVANT® Earns Its Place

Where usage figures run out, and a neutral reading moves the decision forward.



Michiel R. de Boer
Chief Architect, KAIIVANT®

October 2026

Images in this paper are AI-generated.

INTRODUCTION

Why AI decisions get stuck

Ten leadership discussions about AI stall for the same reason: the room reads activity data while debating capability questions. Licence counts, adoption rates and initiative lists show that tools are present and in use. They do not show whether the organisation is getting better at the work that matters.

When the evidence does not reach the question, the discussion turns to opinion. Each executive brings an interpretation. The meeting ends with the decision deferred, or settled by hierarchy rather than evidence.

The Kaivant Score changes what the room reads. It measures nine dimensions across two axes: what the organisation can produce with AI today, and whether the human capability that sustains that output is growing or contracting. The score is non-compensatory, so strength on one axis cannot hide weakness on the other.

This paper sets out the ten discussions where that reading earns its place. Each follows the same pattern: where the discussion stalls, what the room is really asking, what the reading adds, and how the decision moves forward. Start with the one closest to a question already on your agenda.

The ten discussions at a glance

#	DISCUSSION	THE STUCK QUESTION	WHAT THE READING ADDS	INSTRUMENT
01	Justify the AI spend	What is our AI investment returning?	Whether AI multiplies value and whether capability is growing	Kaivant-O
02	Underperforming deployment	Adoption is solid, so where are the gains?	Which of three drag patterns is active	Kaivant-O
03	Reorganisation cover	Is this structural change the right one?	A neutral structural reading in place of a personal debate	Kaivant-O
04	The board is asking	How AI-native are we, really?	A defensible position and a direction of travel	Kaivant-O
05	Capability risk	Output is up, but are people getting better?	The augmentation or substitution pattern, per person	Kaivant-I
06	The AI skills gap	Training is complete, but is the right capability developing?	Whether AI use builds judgment or replaces the work that built it	Kaivant-I
07	Pilot purgatory	Why do pilots succeed but never scale?	The capability gap between the pilot cohort and the wider population	Kaivant-I
08	Leadership bottleneck	Why are decisions slower than the work?	Where strong AI leverage meets constrained decision rights	Kaivant-I
09	Two-speed organisation	The average looks fine, so why does it feel uneven?	The distribution the average conceals	Kaivant-I
10	Agentic AI readiness	Can our people supervise autonomous agents?	Judgment depth, measured before agents act on their own	Kaivant-I

Kaivant-O reads the organisation. Kaivant-I reads each individual privately and shows leaders the pattern across a team.

DISCUSSION 01 OF 10

Justify the AI spend



The board wants to know what the AI investment is returning, and usage figures cannot close that question.

Where the discussion stalls. Management presents licence costs, adoption percentages and active-user counts. These show that tools are present and that people have opened them. They do not show whether the organisation is better at anything that matters, so the board leaves without an answer.

What the board is really asking. The question is not how much was spent or who used it. It is whether the organisation is more capable than before, and whether it can show it. A high adoption rate is consistent with automating low-value tasks while high-value judgment work stays exactly as it was. That is not a return. It is a cost with a visible dashboard.

What the reading adds. The Kaivant Score reads nine dimensions across two axes. Two readings matter most for this discussion.

- **Leverage Trajectory** shows whether AI is applied where it multiplies value, not only where it saves time. High adoption with a declining trajectory means busy tools and flat impact. Low adoption with a rising trajectory means selective deployment that builds capability.
- **Capability Development Velocity** shows whether the people doing the work are more capable than before the deployment. Flat or declining readings point to AI substituting for human judgment rather than augmenting it. That is the answer the board needs, comfortable or not.

How the decision moves. A single score describes the current state. The board wants a trajectory, with a named owner and a named action behind it. In a facilitated session, the leadership team reads the score together, identifies the two or three dimensions pulling it down, and agrees the smallest first move for each. The next board conversation opens on the change between two readings: what moved, and why.

QUESTIONS FOR THE ROOM

Which of our AI investments multiply value, and which only save time?

Are our people more capable than they were before the deployment?

Who owns the next move, and when do we measure again?

DISCUSSION 02 OF 10

Underperforming deployment



The tools are deployed and people use them, but the projected gains have not followed and nobody can say where the gap sits.

Where the discussion stalls. Six months after a major deployment, adoption is on target and nobody avoids the tools. Yet the modelled productivity gains have not appeared, and some teams have become slower. More training is commissioned and change management is reviewed, but the gap remains. The organisation has been measuring adoption, not integration.

What the room is really asking. The tools are not defective and the decision to deploy was not wrong. Something upstream of deployment is not working. The real question is where the drag sits, and which response addresses it.

What the reading adds. Three patterns explain most underperforming deployments. They often run together, and the score shows which combination is active in this organisation.

PATTERN	WHAT IS HAPPENING	WHAT THE SCORE SHOWS
Coordination overhead	Time saved by automation is spent checking, re-running and correcting AI output, and routing work around failures	Coordination Efficiency falls instead of rising
Leverage mismatch	AI went where it was easiest to implement, not where it matters most, so teams are faster at tasks that were never the bottleneck	Leverage Trajectory is flat or declining
Capability substitution	People rely on AI for work they used to do themselves, while the outputs still look fine	Capability Development Velocity reads lower than before deployment

How the decision moves. The score does not presume the cause. It shows where the weight sits across nine dimensions, and the pattern is usually recognisable once the team sees it. A neutral reading replaces an accusation that people would otherwise defend against. A two-hour facilitated session is typically enough to identify the active pattern, agree which dimensions to address and set the first move for each.

QUESTIONS FOR THE ROOM

How much of the time we saved is now spent managing AI output?

Did we deploy where the bottleneck was, or where deployment was easiest?

Would our teams perform as well today without the tools?

DISCUSSION 03 OF 10

Reorganisation cover



A structural change is on the table and the debate has turned personal; a neutral reading turns it back into a structural question.

Where the discussion stalls. A leadership team is weighing a structural change: combining functions, removing a layer, or reorganising a team around different work. The sponsor has a rationale, and several people in the room have another. Without shared evidence, the debate becomes about who is right rather than what the organisation needs. The outcome can feel like a power move rather than a response to a real condition.

What the room is really asking. Beneath the debate sits a shared question. Is the current structure holding back the way work is now done with AI, and if so, where?

What the reading adds. Three Kaivant-O dimensions speak directly to a reorganisation decision. Low readings on these describe a structure, not a person.

- **Coordination Efficiency:** where meetings, approvals, rework loops and handoffs absorb capacity that should create value.
- **Decision Velocity:** where decisions are escalated that could be made locally, and where escalation adds delay rather than quality.
- **Autonomous Workflow Percentage:** how much high-value work flows without unnecessary friction, versus work still routed through structures built for an earlier way of working.

When all three show drag, the room has the structural rationale it needs. Nobody has to make a political argument.

How the decision moves. The team reads the score together before anyone names the change. The facilitator keeps the discussion on what the instrument shows, not on who built the current structure. By the end, the room agrees the structural condition and names two or three changes to address it. A decision that comes from shared evidence is one the whole room can own.

QUESTIONS FOR THE ROOM

What does the evidence say about our structure, before we discuss options?

Which decisions travel upward by habit rather than need?

How will the next reading show that the change closed the gap?

DISCUSSION 04 OF 10

The board is asking



The board or investors want to know whether the organisation is truly AI-native, and investment figures and initiative lists do not answer that.

Where the discussion stalls. A board agenda item, an investor review or a governance query asks how AI-native the organisation is, and whether it is ahead or behind. The reply usually arrives as spend, adoption rates and a list of initiatives under way. These describe activity. They do not describe position.

What the board is really asking. The question has two parts, and neither is answered by a spend figure or an initiative count.

- **The comparative question:** relative to where we need to be, are we ahead or behind? Two organisations with identical adoption rates can sit at opposite ends of a capability trajectory.
- **The durability question:** is our AI integration building something that compounds, or something that depletes? Strong AI activity can quietly erode the human capability behind performance, and no dashboard the board sees today shows it.

What the reading adds. The score gives a structured position across nine dimensions of integration quality. It looks past whether tools are present to whether they multiply value, and whether the capability beneath them is growing. On durability, Leverage Trajectory and Capability Development Velocity show whether integration compounds or depletes. High adoption with flat capability development is an expensive substitution programme. The result is a reading the board can defend: where AI augments capability, where it substitutes for it, and the move planned this quarter to improve the ratio.

How the decision moves. One reading answers the current-state question. The board-level answer comes from the loop: the score, a facilitated session, agreed moves with named owners and timelines, and a set re-measurement date. The next board update compares two readings rather than presenting a new slide deck.

QUESTIONS FOR THE ROOM

Can we describe our AI position in a way we could defend to investors?

Is our AI activity compounding capability or consuming it?

What will the next reading show, and who is accountable for it?

DISCUSSION 05 OF 10

Capability risk



Your team produces more with AI; whether they are getting better at the work is a separate question, and it can be measured.

Where the discussion stalls. Output is up, but leaders sense a change. Work that once needed real thinking now goes to the AI first. People are less willing to take on problems where the tool cannot lead, and some struggle to reason through an argument unaided. The output metrics say all is well, so the concern has no data to stand on.

What the room is really asking. This is the substitution pattern. AI absorbs not only execution but the work that built judgment and range. Output holds steady or improves, so the decline stays hidden until the tools fail or a novel situation demands unassisted judgment. By the time it shows in output data, it has usually been building for over a year, and the options have narrowed.

What the reading adds. The Kaivant-I reads each person on two axes. Personal Leverage Architecture shows how well AI is built into their workflow. Human Capital Depth shows what is happening to the capability underneath. The key dimension here is Capability Independence and Resilience: whether people can still perform at quality without AI. Strong leverage with falling independence is the substitution signature. Because the score is non-compensatory, strong leverage cannot mask declining capability.

How the decision moves. Each person completes the assessment privately, and individual independence results stay with the individual by design. Leaders receive the team pattern: where the substitution signal appears, and whether their concern is confirmed or contradicted. A facilitated session agrees two or three targeted changes. These are usually structural: reintroducing the kinds of work that build the declining dimensions, without giving up AI leverage that adds value.

The goal is not to use less AI. It is to use AI in ways that compound capability rather than substitute for it.

QUESTIONS FOR THE ROOM

What would happen to quality if our AI tools were unavailable for a week?

Where has AI taken over the work that used to develop our people?

Will we measure this now, or wait until it shows in performance?

DISCUSSION 06 OF 10

The AI skills gap



Your team is upskilling in AI, and training completion cannot tell you whether the right capability is developing.

Where the discussion stalls. The training budget is spent, completion rates are high and adoption is up. Yet the way the team handles hard problems feels different: working through uncertainty, synthesising without a template, defending a position from first principles. The usual answer is that training takes time to show results. That may be true, or the training may have measured the wrong thing.

What the room is really asking. Learning to use AI tools well is a workflow question. Developing the judgment that stays valuable as AI use deepens is a capability question. Most AI training programmes are designed to build and track the first, not the second. Completion rates measure participation, not what is happening to capability beneath the tool use.

What the reading adds. The decisive distinction is augmentation versus substitution.

- **Augmentation:** AI takes over execution, and people invest the freed time in harder, judgment-heavy work. Capability grows alongside tool use.
- **Substitution:** AI also absorbs the developmental work, such as the synthesis and the wrestling with ambiguity. Output looks the same or better while judgment quietly weakens.

Training completion cannot tell these apart. The Kaivant-I can, at the individual level, by reading Personal Leverage Architecture and Human Capital Depth side by side.

How the decision moves. Each individual sees their own full results first. Leaders see the pattern across dimensions: where AI leverage is building, and where the substitution signal is forming. A facilitated session agrees specific, modest structural changes that bring back the developmental work the current pattern removed. Progress is checked against a re-assessment after ninety days.

The goal of the training was never AI literacy alone. It was a team that is more capable because of AI, not one that looks more productive while quietly becoming more dependent on it.

QUESTIONS FOR THE ROOM

Was our training designed to build AI literacy or lasting capability?

Where are people investing the time AI saves them?

What do we expect to see at a ninety-day re-assessment?

DISCUSSION 07 OF 10

Pilot purgatory



Every pilot looks promising and none of them scale; the bottleneck is usually a capability gap, rarely a technical one.

Where the discussion stalls. Pilots are small, staffed with motivated people and closely managed, so they almost always pass their own criteria. Then the expansion stalls, never quite happens, or quietly reverts. The tool stays licensed and the initiative stays nominally active, but the work has not changed. The cycle repeats with the next tool.

What the room is really asking. The usual explanations are weak sponsorship, poor change management or a tool that did not fit. Sometimes they are right, but often they are not. The pilot cohort had something the wider group lacks: an existing pattern of AI use, and the judgment depth to absorb a new tool productively. When the rollout widens, a tool that does not fit how people work is quietly set aside rather than openly resisted.

This is a capability gap, not a motivation gap. That distinction changes what the response needs to be.

What the reading adds. The Kaivant-I measures the two things that decide whether adoption lasts. Personal Leverage Architecture shows whether AI fits how each person really works. Human Capital Depth shows whether people can supervise AI output, spot when it leads them astray and perform where it cannot help. Comparing the pilot cohort with the population you plan to deploy into shows exactly where the gap sits and how wide it is.

How the decision moves. The result is not a reason to delay. It is a map of what must be in place for the rollout to hold. A facilitated session turns the comparison into a sequenced deployment plan: which dimensions to develop first, which groups are ready now, and which design changes close the gap better than another round of adoption training.

Whether the pilot worked is not the right question before scaling. The right question is whether the people you are scaling into have what made the pilot work.

QUESTIONS FOR THE ROOM

What did our pilot participants have that the wider organisation does not?

Are we scaling a tool, or the capability that made it work?

Which groups are ready now, and which need preparation first?

DISCUSSION 08 OF 10

Leadership bottleneck



Your AI-enabled people now move faster than your decision structure, and your best performers feel the friction first.

Where the discussion stalls. The tools work, and some people are much faster, with better output, more options and quicker synthesis. Yet decisions move at the old pace, because the structure above the work has not changed. Sign-off sequences built for a slower pace now create queues in front of people who finished hours ago. The organisation often reads this as slow AI adoption and responds with more training.

What the room is really asking. This is a structural mismatch. AI has raised the capacity of individuals faster than the coordination and decision architecture has been redesigned. The risk goes beyond inefficiency to talent. The people most capable with AI have choices, and a structure that blocks them is not one they will stay in.

What the reading adds. The Kaivant-I shows whether each person's AI leverage turns into real output, or whether friction in the surrounding system absorbs the gain. It also reads decision velocity: whether decisions land at the right level and speed, or whether every meaningful move must travel upward. Strong leverage paired with constrained decision velocity is the mismatch signature. Across a leadership team or department, the pattern shows where the drag sits and how far it runs.

How the decision moves. More training is the wrong answer, because the adoption is already there. A facilitated session maps the friction points: which decision types cause the most drag, which roles have high leverage but little autonomy, and where changes would release value fastest. The output is a concrete set of changes to decision rights and approval sequences, made before the talent problem shows up in attrition.

QUESTIONS FOR THE ROOM

Which decisions travel upward because they must, and which out of habit?

Where are our highest-leverage people waiting on approvals?

Would our best AI users say our structure helps or hinders them?

DISCUSSION 09 OF 10

Two-speed organisation



Half the organisation runs on AI while the other half watches, and the adoption average hides both.

Where the discussion stalls. The adoption rate sits in the range the industry considers normal. Look closer and a few people or teams are transformed, often concentrated in one or two functions. The rest use the tools occasionally and fall further behind each month. Investment, talent and deployment decisions are being made on an average that describes neither group.

What the room is really asking. The average makes the high performers invisible as an asset and the lagging majority invisible as a risk. The gap does not close by itself. People with AI leverage develop faster, and the distance grows each month. Left alone, it becomes a structural divide that normal people programmes struggle to bridge.

What the reading adds. Run across a team, function or organisation, the Kaivant-I produces a distribution across nine dimensions, not an average. It shows who has real AI leverage, which dimensions are strong in which groups, and where the lagging group is stuck and why. The fast group shows strong leverage that keeps building. The slower group shows flat readings, with the widest gap in how much of their work has AI truly built in rather than merely available.

Knowing which dimensions drive the gap tells you which problem you face. A workflow gap, a judgment gap and a structural or incentive gap each need a different response.

How the decision moves. One programme for everyone holds back the fast group and fails the slower one. A facilitated session uses the distribution to design parallel tracks. The fast group usually needs structural support: more decision autonomy and harder challenges that compound their leverage. The slower group usually needs workflow redesign, not more training. The results also give leaders the evidence to make the case internally against a single programme.

The number to manage is not the adoption average. It is the distance between the two groups, and how fast it is widening.

QUESTIONS FOR THE ROOM

What does our adoption distribution look like beneath the average?

How fast is the gap between our two groups widening?

Are we giving both groups the same programme because it is easier?

DISCUSSION 10 OF 10

Agentic AI readiness



You are deploying AI agents that act on their own, and whether your people can supervise them is measurable before launch.

Where the discussion stalls. Earlier AI tools were co-pilots: the person drove and the tool helped. Agentic tools initiate, execute and complete sequences of actions across systems. The human is now a supervisor who monitors, catches errors, overrides when needed and stays accountable. Deployment plans often test whether people can use the tools, not whether they can hold this judgment role.

What the room is really asking. Agents produce outputs that look authoritative. Catching a subtle error requires thinking independently about the problem the agent was solving. The people who rely most on AI assistance are exactly those whose independent judgment is most at risk. Deploying agents into a population with weakening judgment depth is a specific, identifiable risk. The agent will make errors; the question is whether the human layer will catch them.

What the reading adds. The Kaivant-I reads Human Capital Depth: whether people still exercise independent judgment, keep the breadth of thinking that supervision requires, and handle novel situations the AI cannot lead. The critical dimension is Capability Independence and Resilience, whether people can perform well without AI. Strong AI integration with falling independence can work when a person directs a co-pilot. It does not hold when that person must evaluate work they did not help produce. The reading shows which people and teams can supervise reliably, and which need capability rebuilt first.

How the decision moves. A facilitated session turns the results into a deployment recommendation: which groups are ready, which need development first, and what that development must address. The intervention is not a training course. It reintroduces the judgment challenges that supervision will demand, in settings where errors are still recoverable. The goal is to make AI leverage safe to extend, not to reduce it.

When an agent makes a consequential error, the person supervising it is accountable. The time to know whether that person has the judgment depth the role needs is before deployment, not after the first error.

QUESTIONS FOR THE ROOM

Who is accountable when an agent makes a consequential error?


Do the people we plan to make supervisors have the judgment the role needs?

Where can we build that judgment while mistakes are still cheap?

THE METHOD

The loop behind every discussion

Each of the ten discussions ends the same way: not with a score, but with a measured change the board can follow.



The developmental loop. The next board update opens on the change between two readings.

A single reading describes where the organisation stands. The loop turns it into a direction. The leadership team reads the score together, agrees two or three moves with named owners, and sets the date of the next measurement. The next board update then compares two readings rather than presenting a new slide deck.

Bringing the score into the room

- **Start with the discussion already on your agenda.** The score earns its place when it answers a live question, not as a general audit.
- **Match the instrument to the question.** Kaivant-O suits questions about the organisation and its structure (discussions 01 to 04). Kaivant-I suits questions about people and capability (05 to 10).
- **Read before you debate.** Share the reading before anyone proposes a solution, so the discussion starts from shared evidence.
- **Protect individual results.** Kaivant-I results stay private to each person by design. Leaders work with the team pattern.
- **Set the next measurement date before the session ends.** Without it, the score stays a snapshot.

A score is a starting point. An agenda, owned by named people and tested by the next reading, is what moves the organisation forward.



Take the next step

Senior leaders

Free Kaivant-O Leadership Pulse

kaivantscore.com/start

Individuals

Free Kaivant-I test

kaivantscore.com/KI-test

Read the full entries

Each discussion has its own entry on kaivantscore.com.

01	Justify the AI spend	kaivantscore.com/decisions/justify-the-spend
02	Underperforming deployment	kaivantscore.com/decisions/underperforming-deployment
03	Reorganisation cover	kaivantscore.com/decisions/reorganisation-cover
04	The board is asking	kaivantscore.com/decisions/the-board-is-asking
05	Capability risk	kaivantscore.com/decisions/capability-risk
06	The AI skills gap	kaivantscore.com/decisions/ai-skills-gap
07	Pilot purgatory	kaivantscore.com/decisions/pilot-purgatory
08	Leadership bottleneck	kaivantscore.com/decisions/leadership-bottleneck
09	Two-speed organisation	kaivantscore.com/decisions/two-speed-organisation
10	Agentic AI readiness	kaivantscore.com/decisions/agentic-ai-readiness



KAIIVANT® is a registered trademark. · kaivant.org · kaivantscore.com