



From AI Activity to Business Value: A CIO/CTO Discussion Paper

Tech Leaders Connect Dinner co-hosted by BCN

Introduction

AI has moved rapidly from specialist technology discussion to boardroom obsession. It's one of the few technology topics that CEOs, investors, customers and employees all feel they need to have a view on. That creates opportunity, but it also creates noise.

For CIOs/CTOs, the challenge is not simply to “do AI”. Most organisations are already doing something with AI, whether formally or informally. The real challenge is to move from scattered experimentation, shadow usage and executive enthusiasm into measured, governed, useful adoption that improves business outcomes.

The discussion at dinner revealed a clear tension. On one hand, AI can remove drudgery, improve productivity, help people work smarter and create entirely new ways of delivering service. On the other hand, it can easily become an expensive distraction: a collection of pilots, tools and board-level proclamations with no clear link to business value.

The conclusion was not that organisations should slow down or avoid AI. It was that CIOs and CTOs need to help their organisations treat AI as part of business change, not as a technology fashion statement.

Moving AI from Interesting Activity to Measurable Business Value

The most important question for any AI initiative is not “where can we use AI?” but “what business outcome(s) are we trying to improve?”

That sounds obvious, but it is often where organisations go wrong. AI is frequently approached as an answer before the question has been properly understood. Leaders

see the technology, feel the pressure to respond, and start looking for places to apply it. This reverses the normal logic of business improvement.

The better starting point is much more familiar: What do people need to do their jobs better?

That question brings AI back into the world of requirements, process improvement and business change. It forces the organisation to understand what is currently difficult, slow, manual, risky or frustrating. AI may then become part of the answer, but it is not assumed to be the answer.

AI is not just a technology issue

One of the strongest themes from the discussion was that AI should not be viewed as “a tech thing”. It may be delivered through technology, but the opportunities, risks and benefits sit across the business.

AI can apply to almost any role where people need to work faster, make better decisions, summarise information, automate repetitive work, improve customer interaction or reduce administrative burden. In many organisations, there is plenty to improve because existing procedures and ways of working are often inefficient.

Thus, AI adoption cannot be owned by IT or the CIO/CTO alone. The CIO/CTO needs to be deeply involved, and may lead elements of the work, but accountability for value must sit with the business. If AI is being used to improve sales, customer service, finance, operations or supply chain activity, the relevant business leaders must be responsible for the value created.

The CIO/CTO's role is to ensure the organisation asks the right questions, chooses appropriate tools, manages the risks, understands the data, and creates the conditions for safe adoption.

The board must define value

There was a strong view that the board needs to engage properly with AI. There's no need for every board member to understand the technical detail, but because AI affects risk, productivity, investment, customer experience, people, operating model and reputation it must be a constant area for review and discussion.

The board needs to agree what value means. For instance, is the organisation seeking cost reduction, faster service, better decision-making, higher quality, improved customer experience, reduced risk, greater scale, or new revenue opportunities? Without that clarity, AI activity becomes difficult to prioritise and almost impossible to measure.

This is particularly important because CEOs are under pressure to show that their organisation is using AI. AI has become one of the first technologies that some leaders

feel they must tell customers, shareholders and investors they are adopting, regardless of whether it is yet making a measurable difference.

That pressure can lead to poor decision-making. CIOs/CTOs need the courage to push back when AI is being treated as an end in itself.

Productivity improvement is still valuable

Not every AI use case needs to be revolutionary. Much of the immediate value may come from helping people complete ordinary work more effectively.

Examples might include:

- reducing repetitive administrative tasks;
- improving document drafting and summarisation;
- helping teams interrogate large bodies of internal knowledge;
- supporting customer service agents;
- improving tender responses and bid quality;
- speeding up research, analysis or reporting;
- identifying patterns in operational data;
- assisting with quality control or compliance checks.

These may not sound transformational individually, but they can make work better, faster and less frustrating. In many cases, AI may delay future recruitment rather than remove existing jobs. It can reduce the time that people spend on low-value work enabling them to spend more time on high-value work such as judgement, relationships, decision-making and improvement.

However, this benefit only appears if the implementation is thoughtful. If employees feel AI is being imposed on them or used mainly as a threat to their roles, adoption will be resisted. If they are engaged properly, AI can be positioned as removing drudgery and improving quality of work.

From centre of excellence to enablement centre

The discussion challenged the idea of AI sitting in a traditional “centre of excellence” that waits for the business to come to it. A more useful model may be an AI enablement function that actively goes into the business, works with teams, understands their work, and helps them find safe and valuable uses of the technology.

This distinction matters. AI value will not be found only in large strategic programmes. It will often emerge from close engagement with people doing the work; they know where

the friction is, they know which processes are painful and they know where information is hard to find or where repetitive tasks consume time.

The role of AI specialists should be to help those teams explore the art of the possible while staying within sensible guardrails.

Why AI Initiatives Get Stuck in Pilot Mode

Most organisations are now running AI pilots. Few are turning those pilots into sustained operational value.

The discussion identified a familiar pattern: enthusiasm leads to experimentation; experimentation creates interesting demonstrations; but then work stalls before it changes the way the organisation actually operates.

This is not unique to AI. Many technology initiatives fail at the point where experimentation must become business change. However, AI makes the problem more visible because the technology is moving quickly, expectations are high, and the perceived pressure to act is intense.

AI as the hammer looking for nails

A recurring concern was that AI can become a hammer in search of nails. When executives, teams or suppliers become excited by the technology, every problem starts to look like an AI problem.

That is dangerous. Business bottlenecks need to be properly understood before solutions are chosen. Sometimes the right answer is process redesign, better data, clearer ownership, improved training, system integration, policy change or automation that does not require AI at all.

If the underlying process is poor, adding AI may simply automate confusion. The organisation may end up with a faster version of a flawed process rather than a better business outcome.

CIOs and CTOs should therefore insist that AI initiatives pass the same basic tests as any other business improvement activity:

- What problem are we solving?
- Who owns the problem?
- What outcome do we expect?
- How will we measure improvement?
- What process, data or behaviour needs to change?

- Why is AI the right tool?
- What risks are created?
- What happens if the pilot works?

The final question is often missed; a pilot should not begin unless there is at least a plausible path to operational adoption.

The cultural fit matters

AI pilots also stall when they do not fit the culture, operating model or risk profile of the organisation.

Some AI tools may simply be inappropriate for a business because of security, regulatory or governance requirements. Others may be technically impressive but unacceptable to employees, customers or professional standards. In regulated environments, careless adoption can create significant consequences for senior leaders.

The organisation needs to understand what is happening at ground level. Shadow AI is already a major issue in many businesses. Employees may be using public tools to summarise documents, draft responses, analyse information or solve problems without formal approval. This can create risk, but it also reveals demand.

The right response is not simply to ban everything. It is to understand what people are trying to achieve and then provide safe ways to do it.

The missing bridge from pilot to operation

AI pilots often remain isolated because the organisation has not done the hard work needed to operationalise them. A demonstration can be built quickly, but production use requires much more:

- integration with existing systems;
- data governance;
- security review;
- user training;
- process change;
- support model;
- monitoring and audit;
- clear accountability;
- cost management;

- human oversight;
- measurement of benefits.

This is where many pilots fail. The technology may work in a narrow setting, but the organisation has not created the conditions for it to work safely and reliably in day-to-day operations nor the budget to cover off the requirements of a productised system.

This is also where CIOs and CTOs need to be pragmatic. Not every pilot deserves to continue. Some will produce only “AI slop”: outputs that appear plausible but are not accurate, useful or safe. Others will reveal that the real issue lies in data quality, process design or unclear ownership.

Knowing when to stop

Leaders need to be willing to kill AI pilots when they are not producing value. Continuing to experiment can feel like progress, but it may simply consume time, money and credibility.

An AI pilot should be stopped or redesigned when:

- the business problem is not clearly defined;
- no accountable business owner exists;
- benefits cannot be measured;
- the output is not trusted by experts;
- risks cannot be adequately controlled;
- the cost of operationalising it outweighs the benefit;
- adoption would require process change that the business is unwilling to make;
- the use case is better solved without AI;
- the pilot has become an experiment in technology rather than business improvement.

Killing a pilot should not be seen as failure. It is a sign of governance maturity. The point of experimentation is to learn quickly, not to keep every idea alive indefinitely.

Human expertise remains essential

A strong concern was the risk of over-reliance on AI and the erosion of human capability. AI can support experts, but it cannot remove the need for expertise.

Experts are needed to judge whether AI output is correct, useful and appropriate. They are needed to spot errors, bias, hallucination, poor reasoning and missing context. They are also needed to understand what the organisation should do with the output.

There is also a longer-term workforce issue. If AI performs tasks that junior employees would previously have used to learn their trade, organisations may create a capability gap. People still need opportunities to develop judgement, pattern recognition and professional skill.

The question is not simply how AI can improve productivity today, but how the organisation maintains expertise for the future.

Adopting AI Safely Without Killing Momentum

The governance challenge is to create enough control to protect the organisation without forcing every AI use case through a slow, traditional project model.

Too little governance leads to shadow AI, data leakage, reputational damage, poor decisions and regulatory risk. Too much governance kills momentum and pushes experimentation underground.

Utilising a single heavy policy and process is usually too heavy handed with AI. Using a risk-based model that distinguishes between low-risk productivity use and high-risk decision-making or automation is usually more effective.

The board must agree the boundaries

One practical recommendation from the discussion was to hold a board-level conversation specifically about AI boundaries.

The board should discuss questions such as:

- What are we comfortable allowing AI to do?
- What decisions must always remain with humans?
- Where do we require human review?
- What data must never be entered into public AI tools?
- Which use cases are encouraged?
- Which use cases are prohibited?
- What level of risk are we prepared to accept?
- Who is accountable when AI-supported decisions cause harm?

This conversation matters because responsibility ultimately sits with the organisation and its directors. If AI output leads to a poor decision, customer harm, discrimination, breach of confidentiality or litigation, the organisation cannot simply blame the tool.

No AI system will be held accountable in the way a director, professional or company can be. There must be human accountability.

Human in the loop is not optional for high-risk use

For important decisions, there must be human oversight. This is particularly clear in areas such as medicine, legal advice, employment decisions, credit, safety, compliance and customer-impacting judgements.

AI may assist, analyse, summarise or recommend, but a suitably qualified human must remain responsible for decisions where harm could occur. That human must also have the competence and authority to challenge the AI output.

This creates a practical issue: human review is only meaningful if the human understands the domain. If a non-expert is rubber-stamping the AI output is not governance, it's theatre!

Black box risk and the need for testing

The discussion also raised concern about black box AI: systems whose internal operation are evolving and thus not fully understood, even by those deploying them.

Where an AI system changes over time, the organisation needs ways to test whether its outputs remain acceptable. One practical approach is to use known test data and compare outputs against expected answers. This requires experts who can interpret the results and identify drift, bias, inaccuracy or unexpected behaviour.

Without that expertise, organisations may not know whether the AI is improving, degrading or simply producing convincing nonsense.

Data, process and people still matter most

AI does not remove the fundamentals of technology leadership. In fact, it makes them more important.

The familiar triad of people, process and technology still applies, with people and process often being more important than the tool itself. Poor data, unclear process ownership, weak governance and low user capability will undermine AI adoption.

Some organisations are using AI as a catalyst to improve data practice. That may be useful, but CIOs and CTOs should be wary of the idea that AI can simply compensate for weak data foundations. AI may expose data problems quickly, but it does not magically resolve them.

Training and graduated access

Training was identified as a critical part of safe adoption. Users need to understand what AI can do, what it cannot do, and how to use it responsibly.

Neither should every user have the same level of access. Organisations may need to provide graduated access based on role, training, data sensitivity and assessed capability. For example:

- general productivity tools for low-risk tasks;
- approved tools for internal document handling;
- restricted tools for sensitive data;
- specialist tools for technical or analytical teams;
- tightly governed environments for high-risk or regulated use cases.

This allows momentum without treating all AI use as equally risky.

Local, private and cloud-based models

Security architecture is also a major consideration. A local or private model may offer greater control, particularly where sensitive data is involved. However, running AI locally can require significant processing power, specialist skills and infrastructure investment. Many organisations have spent years moving away from local hardware and into cloud-first operating models, so this may not be straightforward.

CIOs/CTOs need to help the board understand the trade-offs between convenience, capability, cost, security and control.

Cost will become a board-level issue

AI cost was a repeated concern. Many organisations are still in the early phase of adoption, where tools may appear affordable, bundled or experimental. But the economics are likely to become more challenging.

Large AI providers have invested heavily and will need to recover those costs. As adoption grows, pricing models may change. Usage-based charges, premium tiers, embedded AI costs and supplier dependency could all become significant.

CIOs and CTOs should prepare their boards for this conversation early. AI is not a one-off innovation budget item. It may become a material and recurring operational cost.

Common Blockers to AI Adoption

The discussion identified several blockers that CIOs and CTOs should expect to encounter:

The pace of change: AI is moving faster than traditional investment cycles. Boards and executive teams are struggling to know when to commit, when to wait, and how to avoid buyer's remorse.

Lack of clear business problem: Many AI initiatives are not sufficiently linked to a specific problem or outcome. This makes value difficult to measure and prioritisation almost impossible.

Cost and uncertain return: AI can require investment in tools, integration, data, training, governance and infrastructure. If benefits are unclear, the cost becomes difficult to justify.

Unrealistic expectations: Some leaders expect AI to “just work”. As is clear, though, it needs proper implementation, process design, data quality, governance and user adoption.

Resistance from employees: Employees may worry about job loss, surveillance, deskilling or being “done to”. Adoption requires engagement alongside deployment.

Poor governance or excessive governance: Too little governance creates risk. Too much governance prevents useful experimentation. The right model is risk-based and proportionate.

Weak data foundations: AI depends on data quality, context and access. Organisations with poor data practices will struggle to get reliable outputs.

Over-reliance on AI: Departments and individuals may start relying on AI outputs without sufficient checking or expertise. This creates operational and reputational risk.

Supplier dependency and platform risk: There is concern about future pricing, changing terms, advertising models, data use and vendor lock-in. The old principle still applies: if a service is free, the user or their data may be the product.

Strategic confusion: Random board-level proclamations about AI can create distraction. CIOs and CTOs need to convert enthusiasm into strategy, priorities and sensible execution.

Practical Actions for CIOs and CTOs

The CIO/CTO has a central role in turning AI from noise into value. The following actions emerged from the discussion.

Reframe AI as business change: Do not allow AI to be treated as a standalone technology initiative. Anchor every AI conversation in business outcomes, process improvement and measurable value.

Establish board-level principles: Facilitate a board discussion on AI appetite, risk, accountability, human oversight and prohibited uses. Create clear principles before the organisation scales adoption.

Map current AI usage: Find out where AI is already being used. Shadow AI is not just a risk; it is also a signal of unmet demand.

Create an AI enablement function: Move beyond a passive centre of excellence. Build a capability that actively works with business teams to identify opportunities, provide guidance and support safe adoption.

Prioritise use cases by value and risk: Assess opportunities based on business value, feasibility, data readiness, risk, cost and ease of adoption. Not all AI ideas deserve equal attention.

Build governance by risk tier: Create lightweight rules for low-risk use and stronger controls for high-risk use. Avoid forcing every experiment through the same process.

Require a route to operational adoption: Do not approve pilots that have no plausible route to live use. Every pilot should define ownership, success measures, operational requirements and decision points.

Protect human expertise: Ensure AI supports experts rather than replacing judgement. Consider how junior staff will continue to learn if AI automates parts of their development path.

Invest in user training: Train people not only on how to use AI tools, but also on their limitations, risks, data responsibilities and need for human review.

Prepare the cost conversation: Help the board understand that AI cost may rise and that supplier economics may change. Build financial awareness into AI strategy from the start.

Conclusion

AI is a significant addition to businesses, but it should never be a project, a toolset or a boardroom slogan. It is an enabling technology that can improve business performance when applied to real problems with clear ownership, good data, appropriate governance and engaged people.

For CIOs/CTOs, the opportunity is significant. AI can help organisations remove low-value work, improve decision-making, serve customers better, and redesign processes that have been inefficient for years.

But the risks are equally real. Poorly governed AI can create reputational damage, legal exposure, wasted investment, employee resistance, degraded expertise and strategic confusion.

The organisations that succeed will not necessarily be those that move fastest or talk loudest about AI. They will be those that ask better questions:

- What business outcome are we trying to improve?
- Who owns the value?
- What risks are we willing to take?
- Where must humans remain accountable?
- What data, process and capability do we need?
- How will we know whether this is working?
- When should we stop?

For CIOs/CTOs, the task is to bring clarity to the hype. AI should be encouraged, explored and used, but it should be used as part of disciplined business leadership, not as a substitute for it.

CIOs/CTOs have always recognised that technology should never be the end goal and the same applies with AI. The Board's goal should instead be about creating measurable business value, safely and sustainably.