IT | PARTNERS

AI agents, and the boardroom.

PRESENTED BY
Andrew Johnson

“

Most boards talk about
AI like it’s an IT initiative.
It’s not. It’s a strategy, governance,
and workflow question.

Whose current best use case is “summarise this board paper?”

AN AI MATURITY LADDER

Three levels of AI: only one of them actually changes how you how you govern.

Most organisations stop at the bottom rung. The value compounds as you climb.

PRIMARY

“Summarise this board papers.”

Useful but low value. This is where many organisations stop.

Condense reports
Rewrite content
Generate meeting summaries
Create action lists

INTERMEDIATE

“Help me understand what matters.”

AI starts becoming a governance support capability not just a productivity tool.

Highlight inconsistencies across papers
Compare numbers across reports
Identify missing information
Separate governance from operational issues
Flag risks and dependencies
Detect contradictions between reports

ADVANCED

“Challenge our thinking.”

The real value is not AI generating content, it’s AI improving the quality of organisational thinking.

Connect recommendations to strategic priorities
Assess alignment to risk appetite
Compare decisions against strategic rules
Surface unintended consequences
Highlight gaps between strategy and execution
Stress test assumptions in board papers

Sick of AI waffle?

A RULE FOR THE NEXT TWELVE MONTHS

~~More.~~ Less.

No one is paid per word. Good doesn’t look like being spammed with twenty pages of AI generated material. It looks like one page that earns the room’s time.

AI should never write the report end to end.

~~01~~	Twenty pages of AI generated prose dressed up as analysis.
~~02~~	A polished narrative no human actually reasoned through.
03	**A one page argument, written by a human, sharpened by AI.**
04	**AI output sits in the appendices, in support of the argument, never instead of it.**

If you can’t say it on one page, you don’t understand it yet.

IF YOU ONLY REMEMBER ONE THING

The value isn't ~~AI generating content~~.
It's AI improving the quality of organisational thinking.

Like…. Automating Low Urgency Critical Tasks

The Learning Curve & Chasm

Not only has the chasm widened, but a fundamental step change has also been introduced, one where many now find themselves stuck

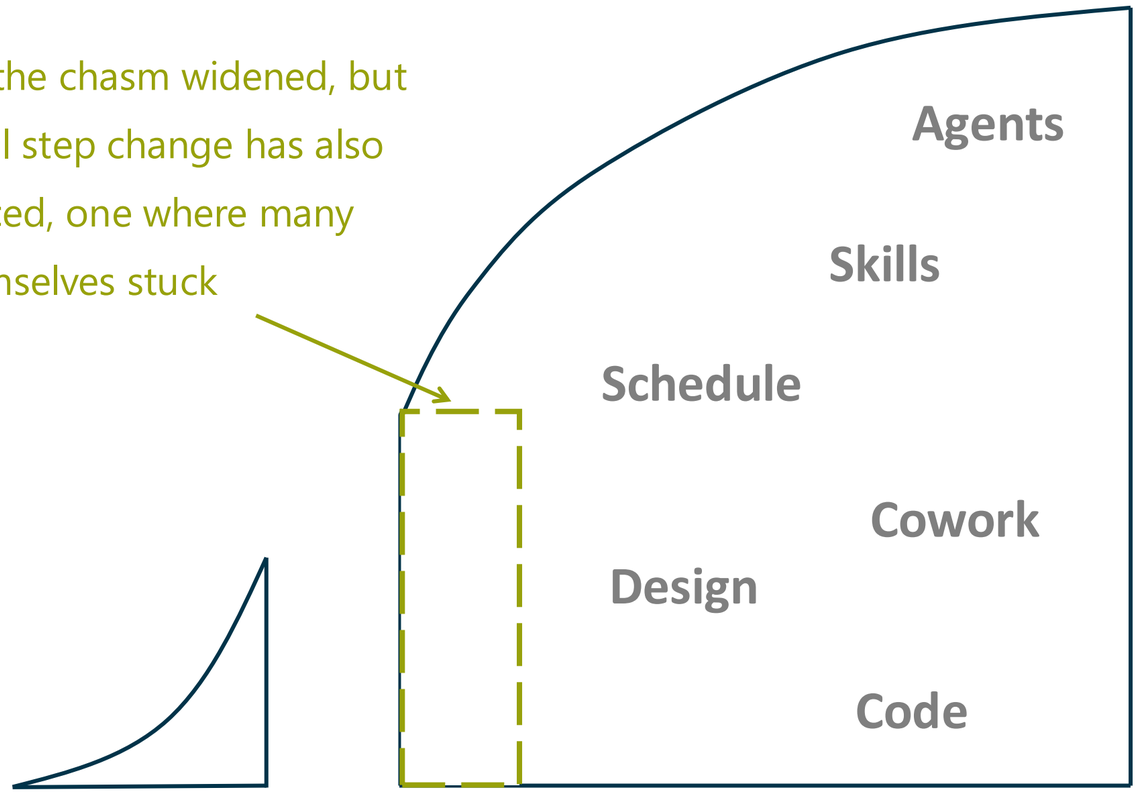

Many of these are advancements from the last 8 weeks alone.

‘Doing AI’ isn’t a strategy.

THE OPPOSITE TEST

If saying the opposite sounds sounds dumb, you don’t have have a strategy.

Strategy is choice. Choosing where to play and how to win. ‘We are going to do AI’ doesn’t pass the test because no one is going to stand up and say the opposite.

Real strategy connects AI to where we play, how we win, and how we measure it through business strategy, operating model, data governance, and tech enablement.

STATEMENT	“We are going to do AI.”
OPPOSITE	“We are not going to do AI.”
VERDICT	Both sound dumb. It’s not a strategy, it’s a buzzword.

STATEMENT	“We will use AI agents to remove 30% of admin from our top three workflows by Q4.”
OPPOSITE	“We’ll absorb that admin manually and stay slower than competitors.”
VERDICT	Choices, trade offs, measurement. Now it’s a strategy.

IT | PARTNERS

WHAT ADOPTION LOOKS LIKE: SOME BASIC EXAMPLES FROM IT PARTNERS

Not chatbots. Agents quietly embedded in everyday workflows.

01

Email Triage

Sort, draft, escalate. Routine correspondence handled before the inbox owner opens it, with humans reviewing what matters.

"The mundane disappears."

02

CV & Team Fit analysis

Compare candidates against the people already in the team, as well as a matrix based on the scope and design of the role.

"Who do we already have? Who's missing?"

03

Industry Change Monitor

Continuously scans the external environment for changes that could impact operations, risk, compliance, vendors, or clients. Surfacing what matters before it becomes urgent.

"What aren't we aware of?"

04

Operational Automation

Time sheet reconciliation, ticket classification, exception flagging. Quiet leverage on the workflows nobody enjoys.

"Less re-work. Same headcount. Higher quality."

Some more sophisticated examples

AGENT HARNESS

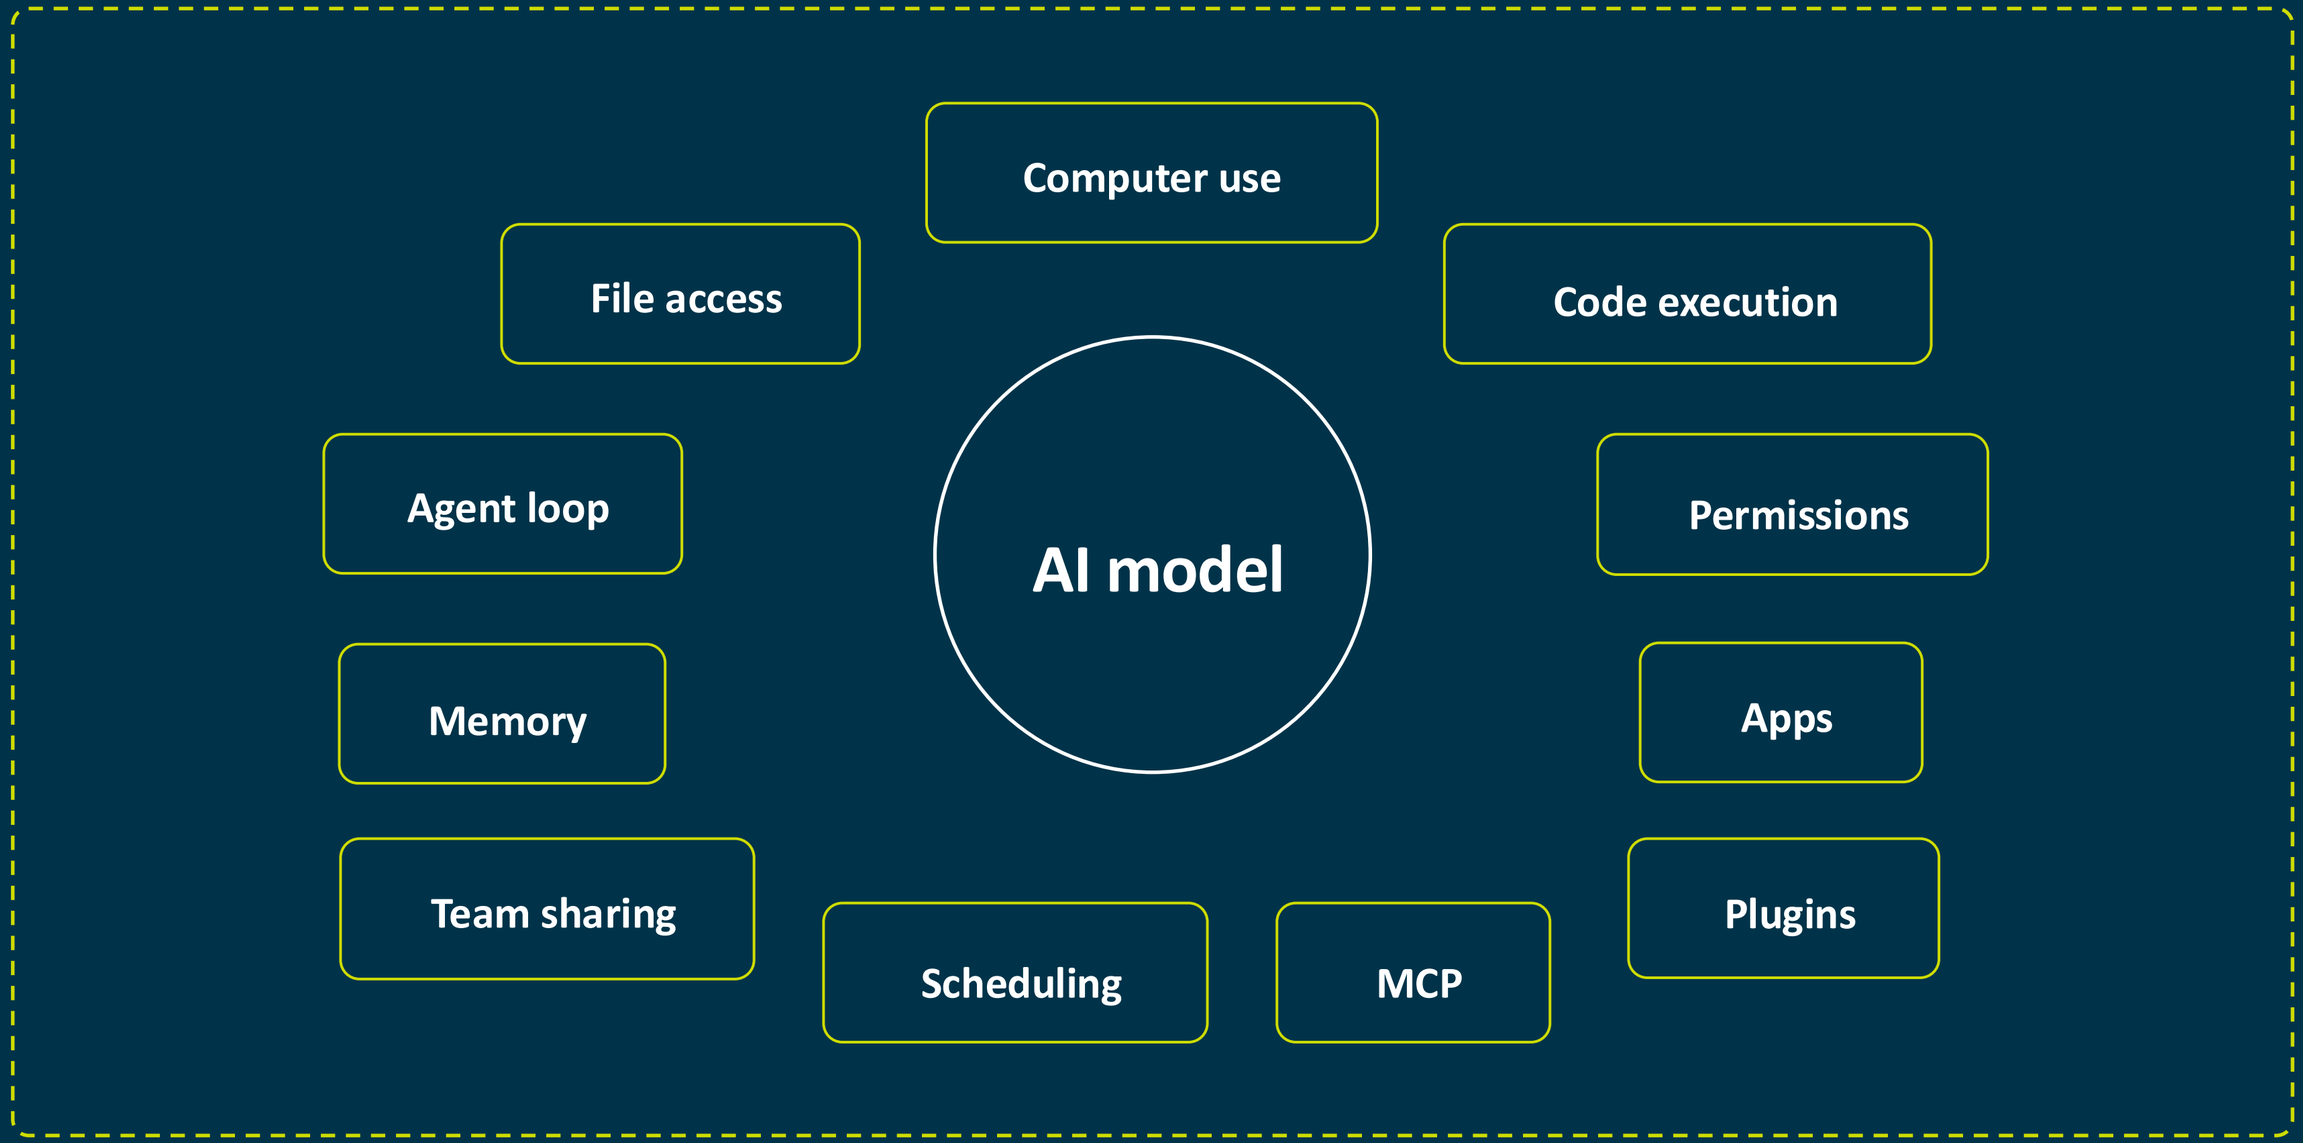

IT | PARTNERS

AGENT HARNESS

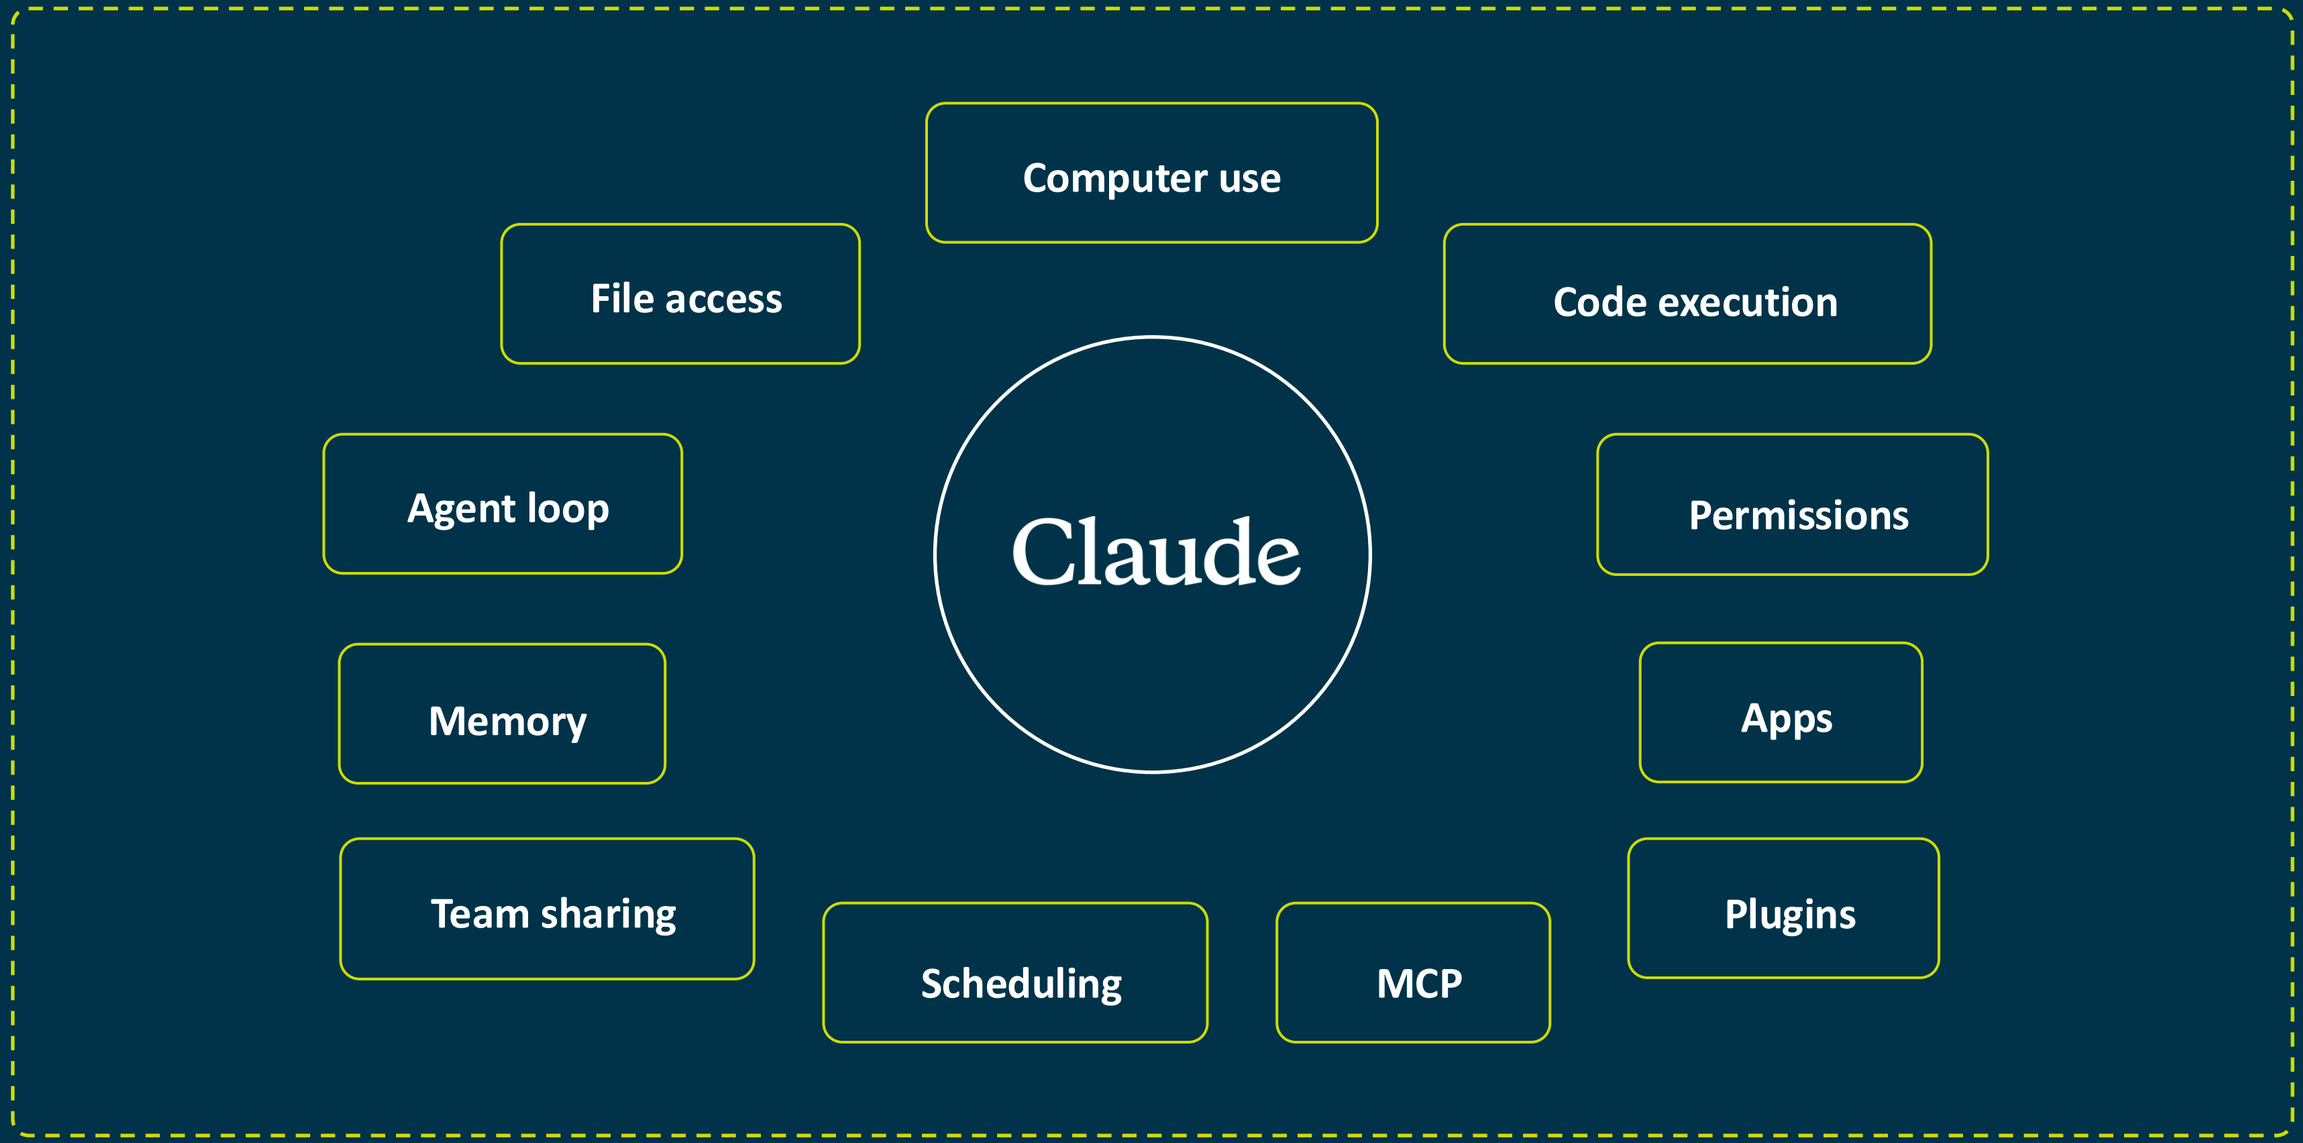

IT | PARTNERS

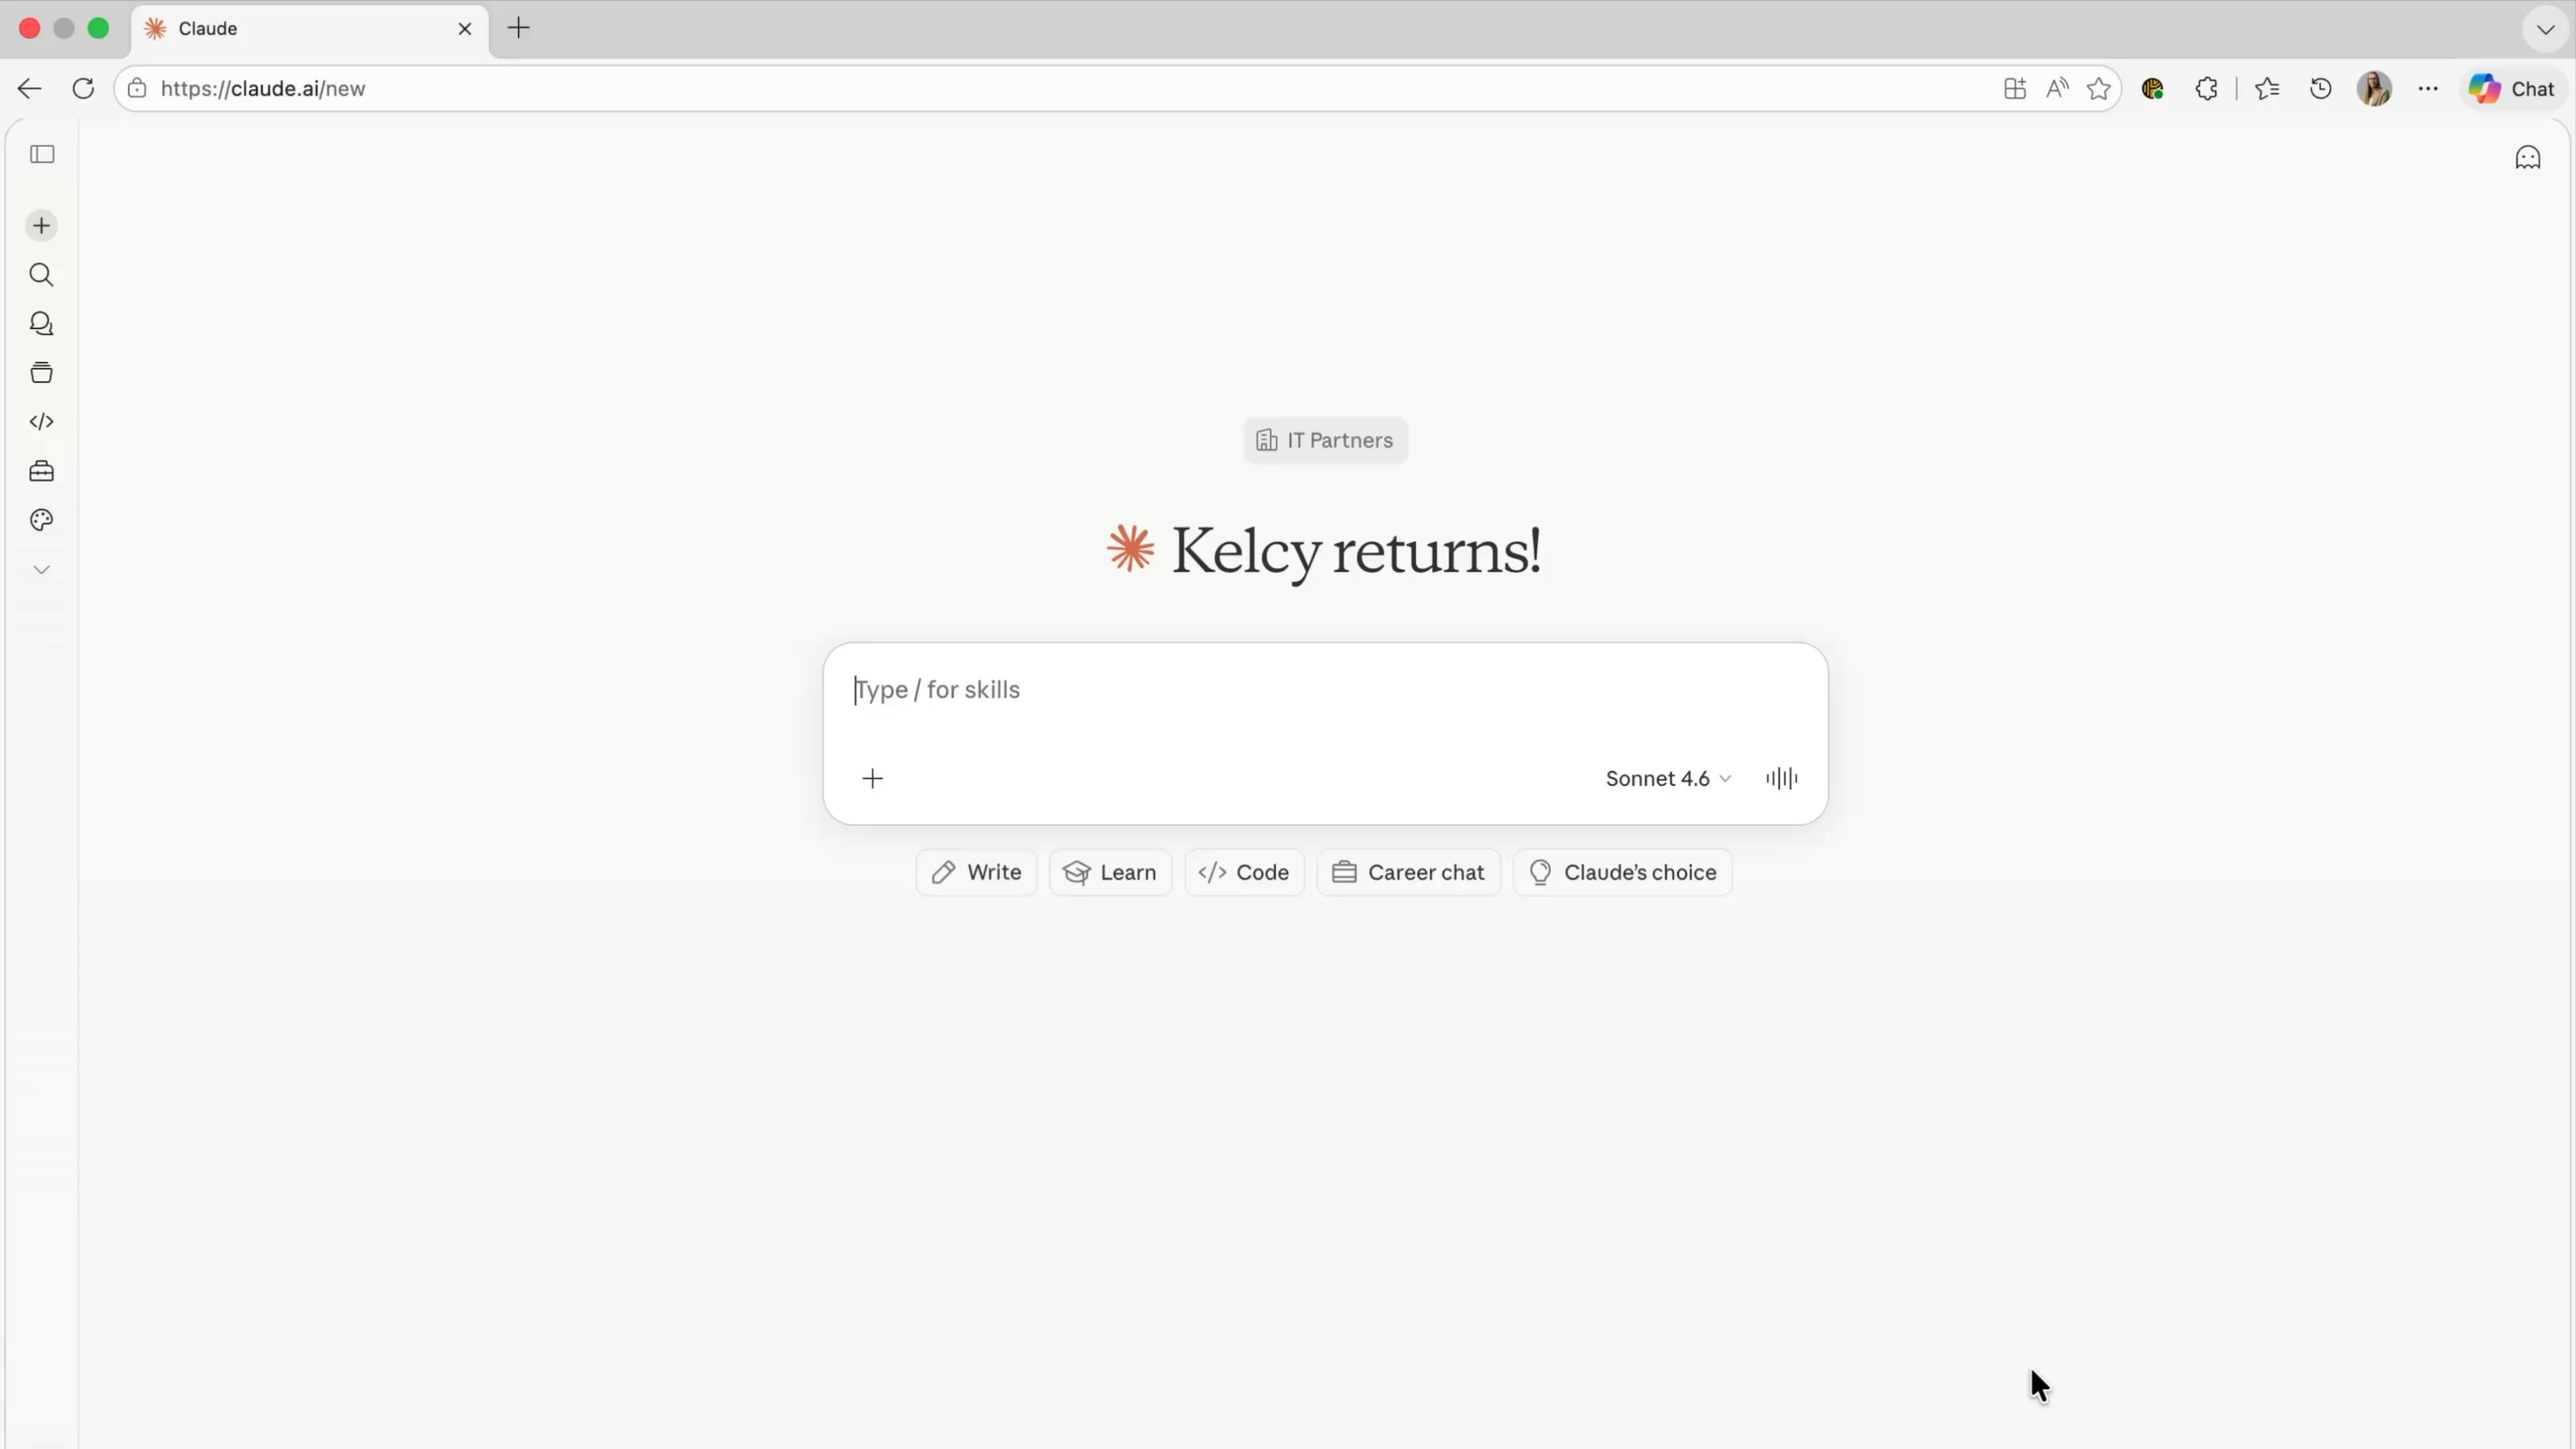
Claude
https://claude.ai/new
Chat
IT Partners
Kelcy returns!
Type / for skills
Sonnet 4.6
Write
Learn
Code
Career chat
Claude's choice

Why did I share this?

Failing to adopt AI is a strategic risk, failing to govern AI is a prudential risk.

AI Governance is now mainstream governance.

So how do we balance the adoption cyber risk?

With every use case we must be considering…

ADOPTION RISK

Most AI projects don’t fail because of the AI.

60%

Most large IT change initiatives fall short of their intended outcomes - the research puts it somewhere between 60 and 70%

They fail because of sequencing, change management, and capability gaps, the same way most tech projects have always failed.

01

Weak sequencing

Project management is only 30% of it. You also need the right people to sequence it, live it, breathe it, and change manage it. New software with no investment in capability is a recipe for waste.

02

No ‘what’s in it for me?’

Communicating change is simple in principle: *What’s in it for me? What’s in it for us?* When was the last time you asked those questions in writing and documented the answer?

03

No honest scoreboard

The discipline most organisations skip: did we achieve what we set out to? Why or why not? If you can’t answer that on a single page, the next investment won’t go any better.

In each use case, where does IT sit?

As agents take on more, the question isn't *can* we automate, it's *should we*, and *where does the human stay*. Three modes. The risk profile is wildly different.

MODE 01 · LOWEST RISK

Human in the loop.

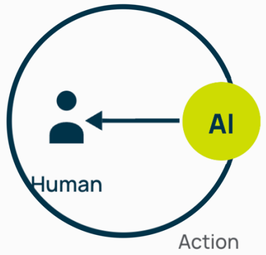

Every material output is reviewed and approved by a person before it ships. AI drafts, humans decide.

Use for: *Anything material: board papers, customer commitments, hiring, financials.*

MODE 02 · MEDIUM RISK

Human on the loop.

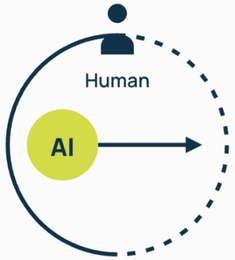

AI acts, a human monitors and can intervene. Speed is real, but so is the chance of a wrong call going out before anyone notices.

Use for: *Low stake ops, triage, drafting, internal automation with audit trail.*

MODE 03 · HIGHEST RISK

Human out of the loop.

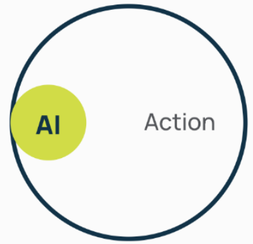

AI decides and acts. No one reviews until something goes wrong, and by then, the decision is in customers, contracts, or the public record.

Default to never: *Unless reversibility, audit, and kill switch are explicit.*

DATA RISK

Your AI is only as safe as the data plumbing behind it.

3x

SURFACE AREA FOR A DATA LEAK THE MOMENT AI TOOLS ENTER THE WORKFLOW

Most AI incidents don't look like sci-fi. They look like a contract pasted into a free chatbot, an over permissioned mailbox, or a model with read access it should never have had.

01

Controls & access

An AI agent inherits the permissions of the account it runs as. If everyone can see everything, so can the model and so can anyone who prompts it. Least privilege isn't paperwork; it's the blast radius.

Ask: *What can our AI tools see that a junior staffer can't?*

02

Shadow AI

Staff are already using free tools with or without permission. Customer data, board papers, and IP are walking out through browser tabs. If you haven't named the sanctioned tools, you've effectively sanctioned all of them.

Ask: *Which tools are on the white list, and how do staff know?*

03

Safe tools

A sanctioned tool with enterprise terms, no training on your data, audit logs, and SSO is a different animal to the free version. Pay the small price now or pay the disclosure cost later.

Ask: *Do our tools have enterprise terms, SSO, and audit logs?*

Board ask: Who owns 'data plumbing' for AI in this organisation and when did we last test it?

CYBER RISK

AI is rewriting attacker economics . Defenders have to catch up.

Anthropic's Project Glasswing has already used a frontier model to surface thousands of zero-day vulnerabilities. The same capability is available to the other side and arriving in headlines every week.

RECENT HEADLINES · LAST 30 DAYS

86%	**of phishing in 2026 is AI-driven.** KnowBe4 Phishing Threat Trends, 2026
APR 28	**AI coding agent wipes production database during weekend deploy.** Mid-market SaaS · reported recovery: 38 hours
APR 11	**Misconfigured AI assistant exposes 1.2M customer records via public API.** Disclosed via NZ Privacy Commissioner Commissioner
MAR 22	**Voice-cloned CFO authorises NZ$2.4M transfer.** Five-second deepfake on Teams call

01

Basics are now critical

Patching, reboots, MFA, backups. AI shrinks the window from disclosure to exploitation from weeks to hours. The unglamorous list is the difference between a Tuesday and a press release.

Ask: *What's our patch cadence and when did we last prove it?*

02

Assess like an attacker

Annual pen tests on a compliance calendar no longer reflect how attacks happen. Continuous testing and AI-assisted assessment are the new baseline, not a premium upgrade.

Ask: *When did we last have a third party try to break in?*

03

APIs & integrations

Every connector, plugin, and MCP server is a doorway. AI agents inherit whichever access level the integration was wired with often more than anyone reviewed. Map what's exposed before someone else does.

Ask: *Which AI tools have read access to our CRM, finance and email today?*

04

AI on the defence too

Glasswing partners (Microsoft, Google, AWS, CrowdStrike, Cisco) are already using AI to find vulnerabilities before attackers do. The board question is whether your providers are on that curve, or behind it.

Ask: *Are our security partners using AI defensively or still scanning quarterly?*

Board ask: If a zero-day was disclosed at 9am, how many hours would it take us to know we're exposed and to be patched?

FOR DIRECTORS

Seven questions every board should already be asking.

Separate governance from operations. Then make sure human judgement is still in the room where the decision happens.

Q1	**Where** is AI already being used in this organisation, including the bits no one's told us about?
Q2	What **assumptions** are baked into the outputs we're relying on?
Q3	For every material decision: is **human judgement** still embedded or are we rubber rubber stamping?
Q4	How are we handling **data**, what's going in, who can see it, where does it live?
Q5	Who is **accountable** when an AI influenced decision goes wrong?
Q6	Are we measuring **real ROI** or accepting inflated vendor numbers?
Q7	What **capability** are we building in our people, not just our tools?

IT | PARTNERS

Three things to take back to your board.

01

AI governance is now mainstream governance.

Don’t silo it. Same disciplines, same questions, applied to a new surface. If it’s not on the agenda yet, that’s the agenda.

02

Practical experimentation beats theoretical debate. Get help.

Two or three controlled experiments, with measurement, will teach you more than another year of strategy papers.

03

Strategy × people × data × execution wins.

The organisations that get all four aligned will outperform: quietly, compounding, before anyone notices.

IT | PARTNERS

Making each day better through technology