

A BRIEF FOR THE CEO AND THE LEADERSHIP TEAM

Our point of *view*.

*Clients normally brief the agency.
Here we brief you: what it actually
takes to get ROI from AI, now the
easy part is over.*

Executive summary

READ THIS PAGE IF NOTHING ELSE

The first wave of AI is done. You bought the licences, you saw the productivity demos, you watched your people play with the tools. It helped, a little. And then the ROI stalled, because **the licences were never the hard part**. The hard part is that AI runs on how your organisation actually works, and most of that was never written down.

This is not an IT project and it cannot be delegated to one. It is the biggest **organisational design** question of the decade, and it lands on the CEO and the leadership team. AI does not transform your organisation. It reveals it, the undocumented workflows, the broken handovers, the unclear ownership, and then it asks you to rebuild. That rebuilding is leadership work, and it is the work most companies are about to skip.

What we believe

You have to **experience** AI to lead it, not read about it. You need a **why** that goes beyond productivity. And AI will **stress-test your foundations**: your ways of working, your incentives, eventually the deal between employer and employed.

What it takes

Undivided leadership time on your real work. A stance on what stays human. A reason for your people to hand over the knowledge in their heads, **because they will hide it if the exchange feels unfair**. This is culture and design work, done inside the work, not in a classroom.

THE QUESTIONS TO SIT WITH, AS A LEADERSHIP TEAM

- ? What do you actually believe AI is here to do to this company: make it efficient, augment it, or refound it?
- ? Which decisions stay human because of who you are, not because the model is not good enough yet?
- ? If AI removes the effort where judgment is forged, where will your next senior experts come from?
- ? If AI changes the work, what do you owe the people doing it, and can they feel it?
- ? Who can call you out when you break your own rules, and what happens when they do?
- ? What will you never do with AI, even when it is legal, cheap and nobody would notice?

The pages that follow set out the argument in full: the house AI walks into, the three things that change at once, our point of view, and the AI Constitution, the one document that turns a policy nobody believes into a deal your people will actually sign. It closes with a bold idea for making that deal unbreakable, and with who we are.

AI removes the buffer.

For two centuries the mess was survivable. *AI is the first thing that cannot work around it.*

For two centuries, an organisation could carry undocumented knowledge, unclear ownership, informal power and contradictory incentives, and still function. Human thinking was slow and expensive enough to quietly absorb the mess. That slack was a buffer.

AI removes the buffer. The moment you put AI on a process, the first thing that breaks is not the technology, it is a handover nobody ever wrote down. This is why the ROI never arrives on schedule: the workflows are undocumented, the handovers are inconsistent, and the knowledge lives in heads. Licences do not fix that. **Access without workflow change is theatre.**

The mess has an address. *It lives in the attic, the basement, and every floor between.*



THE ATTIC *tacit knowledge*

The know-how in people's heads, never written down. What AI needs most, and hardest to reach.

THE MESSY MIDDLE *broken handovers*

Every floor between: assumptions, quiet workarounds, undocumented handovers. The first thing AI breaks.

THE BASEMENT *organisational shadow*

Unclear ownership, informal power, contradictory incentives. Foundations everyone knows, no one names.

The work, the team, the organisation.

Most programmes fix one of the three. *Then the other two surprise them.*

01 · The work

Workflows get rewired.

What took weeks takes days. Outputs shift from slides to working artefacts. The value of knowing-that collapses, the value of judgment climbs.

02 · The team

Handovers get exposed.

Every assumption, every inconsistency, every quiet workaround becomes visible. Psychological safety stops being a nice-to-have and becomes infrastructure.

03 · The organisation

Foundations get stress-tested.

Structures, incentives, and eventually the employer-employee contract itself. AI is the biggest organisational design question of the decade.

AI changes all three at once. We are challengers, so we name it: the attic and the basement are there, everyone in your building knows they are there, and with a mirror we can show you. **Naming it is the first step to rebuilding it.** That is what organisational architecture means to us: not moving boxes, but redesigning how the work, the judgment and the knowledge actually flow.

Three things we hold to be true.

Everything we build follows from these three.

POV 01 You have to experience AI to understand it.

Reading and talking about AI is knowing *that*. Leading with it demands three other kinds of knowing, how, what matters, and through co-creation, and they only arrive through doing, in a room, on your own work.

POV 02 You need a **WHY** that goes beyond productivity.

Efficiency gains are table stakes. The real why has to be specific to your organisation and central to its strategy, and it cannot be delegated. The original thinking is CEO work: a rallying cry, a manifesto, and eventually a constitution.

POV 03 AI will stress-test your foundations.

AI runs on your tacit knowledge. But why would your people write down the very know-how that makes them hard to replace, so it can be automated? They will not, and they should not, not for free. We advocate skin in the game, on both sides: contribution recorded, recognised, rewarded.

Your why, written down. *With skin in the game.*

A policy tells people what they may not do. *A constitution tells them why they should bother.*

Most AI documents are permission slips written by legal. Nobody reads them, nobody believes them, and they answer none of the questions employees are actually asking: why does this matter, what is in it for me, and what happens to me. The AI Constitution is the CEO's answer to all three. Its substance is a new deal between employer and employed, written in the leadership team's own words around the CEO's personal preamble.

WHAT EVERY COMPETITOR SELLS The AI policy. Written by legal and IT, reviewed by compliance, filed on the intranet. It answers what is forbidden. Defensive by design, dated on arrival, read by almost nobody. It creates not one gram of belief.	WHAT WE FACILITATE The AI <i>Constitution</i>. A vision with red lines. Why this matters here, what stays human, what each level may decide, what the company owes its people in return, and how the ones already using AI in secret come home. If the CEO cannot recite it, it is too long.
---	--

The question nobody else asks. If documenting your tacit knowledge is what makes AI work, who gets rewarded for it? Your people will not trade decades of hard-won know-how for a signature on a promise a future leader can quietly revoke. The serious answer is to make the reward real and durable: recorded, recognised, and set aside so a change of leadership cannot reclaim it, the way companies already escrow retention pay. Skin in the game, on both sides.

The anatomy. One preamble and six short articles, the whole thing recitable.

Preamble THE MANIFESTO · NON-DELEGABLE	The CEO's answer to why anyone should care. Not a strategy statement, a personal position: what I believe this means for us, why I am asking you to engage, and what I promise in return. Written by the CEO, defended by the CEO.
---	---

Article I THE ACCESS COVENANT	You cannot ask people to sense a revolution through a keyhole. If you expect your people to grasp frontier AI, you owe them access to frontier tools. Tool access becomes the employer's obligation, not a perk rationed by IT.
--	--

Article II THE RED LINES	The non-delegables, named. Which decisions and judgments stay human regardless of capability, written by your leadership team, not imported from a template. The protective work of a forty-page policy, in one page.
---	--

Article III

THE DISCRETION
LADDER

Inside the red lines, each level holds explicit discretion to apply context in its own area. Compliance to the frame, freedom within it. This is what makes it livable rather than laminated.

Article IV

THE KNOWLEDGE
COVENANT

The boldest article. The company rewards people for externalising the tacit knowledge in their heads, and makes that reward durable enough to trust. The research is unambiguous: employees respond to AI by hiding what they know, whether or not the employer ever signals displacement. This article answers the benefit question in a form a person can feel. Its engineering is on the next page.

Article V

THE AMNESTY

A constitution that ignores the AI already happening in secret governs a fiction. A standing, no-sanction disclosure process: hidden practices surface, and each becomes a supervised experiment. Not a one-off campaign, a permanent clause. The amnesty is what makes the rest of the document true.

Article VI

THE REVIEW

A living document with a re-signing cadence, reviewed against what has actually changed. A constitution frozen in the year of its signing becomes a policy by its anniversary.

The product is the process, not the paper. A template constitution is worthless by definition. The writing is the intervention: two facilitated sessions with the CEO on the preamble, one drafting day with the intact leadership team on the articles, a public signing, then the cascade, where each level localises its discretion and the amnesty window opens. Roughly three weeks, one full leadership day. In German organisations the works council co-shapes Article IV from the start, as co-author, not obstacle.

The promise that survives its promiser.

Ask people to hand over what makes them hard to replace, and a promise is not enough. *The serious answer is to make it unbreakable.*

This is the future-forward end of the argument, the one for a leadership team ready to think several moves ahead. Article IV asks a great deal: hand the company the tacit knowledge that is your quiet job security. No rational person accepts that trade against a promise the employer can revoke at no cost. And paper promises die three deaths: **succession** (the leader who signed it leaves), **reinterpretation** (the successor honours the letter and guts the spirit), and **opacity** (nobody can verify what was paid against what was earned).

So the bold CEO does something no policy can: makes the reward unbreakable. The reward for each contribution is written into a **smart contract** and notarised on a public **blockchain**, used here as a neutral, tamper-proof ledger, not as money. Four moves:

R.01 · Attest

Record the contribution.

Each accepted contribution, a documented workflow, an encoded heuristic, a colleague trained, is hashed and anchored to the ledger. Content never touches the chain, names never touch the chain. Only a tamper-proof fingerprint that it happened.

R.02 · Escrow

Fund the reward upfront.

The reward budget is committed into the contract in advance and released against attested contributions. **Once earned, irrevocable by design:** no successor, no reorg, no budget freeze reaches money already committed. The upfront funding is the employer's own skin in the game.

R.03 · Govern

Hold the terms neutral.

A neutral administrator holds the escrow and the identity mapping, the same third-party role a retention escrow already uses, so management cannot quietly rewrite the terms once contributions are flowing. Changes follow the constitution's own review clause, in the open.

R.04 · Prove

Prove fairness without exposure.

Modern cryptography lets an employee prove their entitlement, and the company prove the programme's fairness to the outside world, without exposing a single person or trade secret. The knowledge stays inside the firm. Only the proof of fairness leaves the building.

Say it precisely. This is a notary, not a bank. No coin, no token, no tradeable anything, no speculation. Entitlements are non-transferable and payout arrives in euros through payroll. The blockchain acts only as the neutral record that makes the promise impossible to quietly break. The words that belong in the room are **notarised, tamper-proof, trust-minimised**. The words that do not are crypto and Web3.

“Your employees are right not to trust your AI promises. You might not be here to keep them. So make the promise one you couldn't break even if you wanted to.”

A provocation for the CEO office, not a procurement item for next quarter. It is here to open a different conversation about what a fair, durable deal with your people could actually look like, and to signal that the

company means it.

The easy part is over.

Past the licences, past the productivity demos, past the token-maxing. *Now comes the organisational work.*

The first phase of AI rewarded curiosity: buy access, run demos, let people play. It was necessary, and it is finished. What is left is the part no vendor sells and no licence delivers: redesigning how work, judgment and knowledge flow through your organisation, and giving your people a reason to come with you. That is organisational development, and it is slow, human, and unavoidable.

It is also where the returns are. The teams that treat this as a leadership problem, not a tools problem, will compound: what you build this quarter compounds next year, and the year after it is exponential. The teams that keep buying licences and waiting for ROI will watch the gap widen, quarter after quarter, because their calendars look the same as everyone else's.

Seeing your own organisation clearly enough to change it is the leadership work now. The 1-Day Executive Experience is the paid proof point: your leadership team feels the frontier on its own work, in a single day, and leaves with the courage and the business case to go further. It is the door into everything else. Bring us a real problem, and we will show you.

Start the conversation → calendly.com/marcusdrueen/challenger-house-intro-call · challenger-house.com

Who is briefing you.

Three operators, one shared mission: new organisational architectures for a post-AI world. Three decades of systemic organisational development, facilitation and venture building, now turned on the AI transition.



Marcus Druen

FOUNDER · ORGANISATIONAL ARCHITECT

25 years as an OD practitioner, facilitator and coach inside blue chips including Microsoft, Telefónica, Sanofi and E.ON. His craft is making the tacit explicit and the messy middle visible. He reads AI as an organisational design problem, not a tool problem.

hello@marcusdruen.com · +44 7947 484882 · linkedin.com/in/marcusdruen



André Cramer

CHIEF LEARNING & AI EXPERIENCE DESIGNER

25+ years in tech, from Silicon Valley (Yahoo!) to Deutsche Telekom, where he led Strategic Communications and the Think Tank for the CTIO board member. Author of the Plausibility Illusion framework, keynote speaker, and curator of the DRANBLEIBEN newsletter and podcast.

linkedin.com/in/theandrecramer · theandrecramer.substack.com



Brittney Bean

BUILDER & CODER · AGENTIC SYSTEMS

Exited serial founder-operator. She builds and implements AI across SMEs and regulated environments, where security, compliance and data governance are not optional. The safety net: she makes sure what gets built survives procurement, legal review and IT governance.

linkedin.com/in/brittneybean