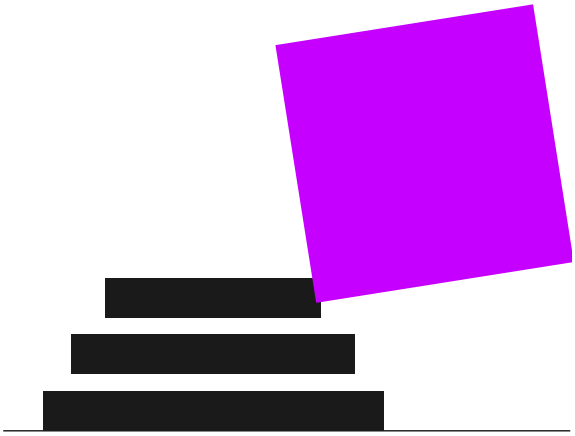

INSIGHT · ORGANIZATION DESIGN FOR AI

Train the Monkey First.

Why your AI operating model is the hard problem you're not solving — and what the 6 % of companies generating real EBIT from AI do differently.

AUTHOR	SERIES	EDITION	PUBLISHED	READING TIME
Maximilian Stein	Pillar 02 · Org Design for AI	Vol. 01 · № 04	May 2026	14 minutes

EXECUTIVE SUMMARY

The infrastructure is built. The organization isn't.

Five findings senior leaders should weigh before approving another quarter of AI capex.

01

Spending is hitting record highs while organizational redesign barely registers.

Global AI spending will hit **$2.52 trillion in 2026** — up 44 % year-over-year. Enterprise budgets allocate 30–40 % to infrastructure and only 5–10 % to training and change management.

02

80 % of AI adopters have not redesigned a single workflow.

Only 21 % of organizations using generative AI have redesigned even some of their workflows. The rest are layering AI on top of pre-existing structures and reporting lines.

03

The 6 % generating real EBIT did one thing differently.

High-performing organizations are **3.6× more likely** to have redesigned their org alongside AI deployment. 55 % of them fundamentally reworked workflows. The other 94 % did not.

04

Operating-model redesign is harder than buying compute — and that's why it gets skipped.

93 % of executives cite culture and organizational readiness as their main AI barrier. The pedestal is fast, tangible, and easy to approve. The monkey is slow, ambiguous, and politically charged.

05

There are six structural domains where redesign is non-negotiable.

Strategy & Value, Structure, Decision Architecture, Process & Workflow, Capabilities & Culture, and System Governance — the six places the monkey hides, and the agenda for the next four quarters.

CONTENTS

What's inside.

01 **The Pedestal Problem** 05
Why $2.5 trillion of AI spending is flowing into the wrong line items.

02 **What Training the Monkey Looks Like** 09
Inside Freeport-McMoRan's organizational redesign at Bagdad.

03 **Why 94 % Keep Building Pedestals** 12
The political economy of organizational redesign.

04 **The Six Places the Monkey Hides** 15
A structural diagnostic for AI operating-model maturity.

05 **What Monday Morning Looks Like** 18
A 90-day sequencing plan for moving from pilots to scaled impact.

06 **About the Author & References** 21
Sources, methodology and how to engage HandsOn.

SECTION

01

The Pedestal Problem.

Astro Teller's thought experiment about training a monkey to recite Shakespeare is a precise diagnostic for the $2.52 trillion enterprise AI economy: nearly all of it is flowing into the pedestal.

The pedestal is fast, tangible, and easy to approve. The monkey is none of those things.

Astro Teller runs X, Alphabet's moonshot lab. He uses a thought experiment to explain how his teams decide where to invest their time: imagine your goal is to train a monkey to recite Shakespeare while standing on a pedestal. You have two tasks — train the monkey and build the pedestal. Where do you start?

His answer: spend zero time on the pedestal. The pedestal is trivially solvable. The monkey is the hard problem. If the monkey can't be trained, the pedestal is a waste of materials and hours.

And yet, as Teller points out, teams gravitate to the pedestal first — because building it creates visible, demonstrable progress. You get something to show leadership. Training the monkey produces nothing but a long list of reasons why the problem is hard.

Global AI spending will hit **$2.52 trillion in 2026** (Gartner). That number is up 44 % year-over-year. Most of it flows into infrastructure — servers, accelerators, data center capacity. The pedestal.

Meanwhile, **80 % of organizations** that have adopted AI have not redesigned a single workflow around it (McKinsey State of AI, 2025). The organizational structures supposed to absorb AI at scale — the roles, the decision rights, the team designs, the workflows — remain untouched. The monkey sits in the corner, ignored.

This piece argues that the operating model is the monkey, and almost every enterprise AI program is building pedestals.

The budget tells the story.

Enterprise AI budgets allocate 30–40 % to infrastructure (compute, platforms, data pipelines), another 30–40 % to model development and API costs, and 5–10 % to training and change management. Organizational redesign — the structural changes required to actually absorb AI into how a company works — barely registers as a line item.

"AI adoption is fundamentally shaped by the readiness of human capital and organizational processes — not merely by financial investment."

JOHN-DAVID LOVELOCK · GARTNER · JANUARY 2026

Figure 01 · Where the AI dollar goes.
Typical enterprise AI budget allocation, % of total program spend.

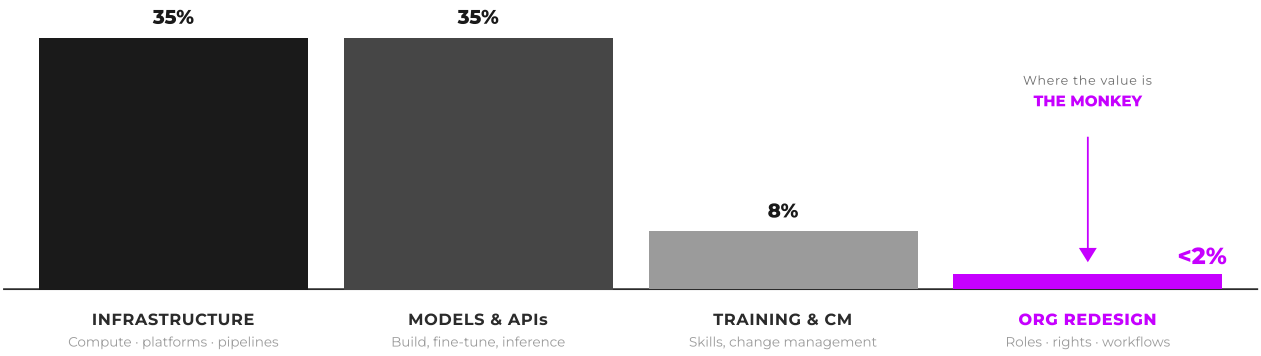

SOURCE · EC-COUNCIL ENTERPRISE AI BUDGET STUDY, 2025; HANDSON ANALYSIS

The readiness isn't there.

A new tool dropped into an old structure is the textbook definition of building the pedestal. The data on workflow redesign is unambiguous — and unflattering.

21%

Of organizations using GenAI have redesigned even *some* workflows.

MCKINSEY STATE OF AI · 2025

80%

Are layering AI on top of pre-existing reporting lines and processes.

MCKINSEY STATE OF AI · 2025

93%

Of executives cite culture and organizational readiness as the primary barrier.

HBR / PYMNTS INTELLIGENCE · 2025

DEFINITION · THE PEDESTAL

A pedestal is anything you build before you have proof the monkey can be trained.

Servers, models, and APIs are pedestals when they are commissioned ahead of decisions about who is authorized to act on AI outputs, which workflows will rely on them, and how the organization will absorb their recommendations. The pedestal is not bad. The pedestal is the easy half. The mistake is mistaking the easy half for progress.

The CTO can present a procurement decision to the board in a single slide. The CHRO explaining why twelve roles need to be restructured, three new team configurations need to be piloted, and decision-making authority needs to shift from function heads to cross-functional leads — that presentation doesn't fit on a slide. It fits in a transformation program that takes quarters, not weeks.

SECTION

02

What training the monkey actually looks like.

At Bagdad, Arizona, Freeport-McMoRan turned a stalled AI deployment into **$350–500 million in projected EBITDA** — not by upgrading models, but by restructuring teams.

Same models. Different operating model. Materially different result.

Freeport-McMoRan, the world's largest publicly traded copper producer, ran into the wall and broke through it. The company deployed AI models to optimize copper processing at its aging Bagdad mine in Arizona. The technology worked — the models could predict and improve ore concentration yields. Scaling across operations failed.

The problem was organizational. AI-generated recommendations sat unused because the people who operated the mills didn't trust outputs from a system they hadn't been part of building. Data scientists and metallurgists spoke different languages. Optimization suggestions that crossed departmental lines had no owner.

The general manager at Bagdad assembled a cross-functional team that pulled from every division the AI initiative would touch: data scientists, metallurgists, mining engineers, and members of Freeport's central data-science group. They adopted agile sprint cycles. They created a product manager role for AI-driven process changes. They moved to quarterly OKR planning across operations.

After the organizational redesign: copper production rose 5 %, throughput exceeded 85,000 metric tons per day — a 10 % increase — and the copper recovery rate improved by a full percentage point.

System-wide implementation is now projected to yield an additional **200 million pounds of copper per year**, worth $350–500 million in EBITDA (McKinsey).

Freeport didn't have better AI. They had a better operating model around the AI. They trained the monkey.

The numbers that the org-redesign produced.

The technology stack at Bagdad didn't change between "before" and "after." What changed was who was authorized to act on AI outputs, which roles owned the workflow, and how cross-functional teams planned together.

Figure 02 · Before and after the redesign at Bagdad.

Operational outcomes following Freeport-McMoRan's organizational redesign around its existing AI models.

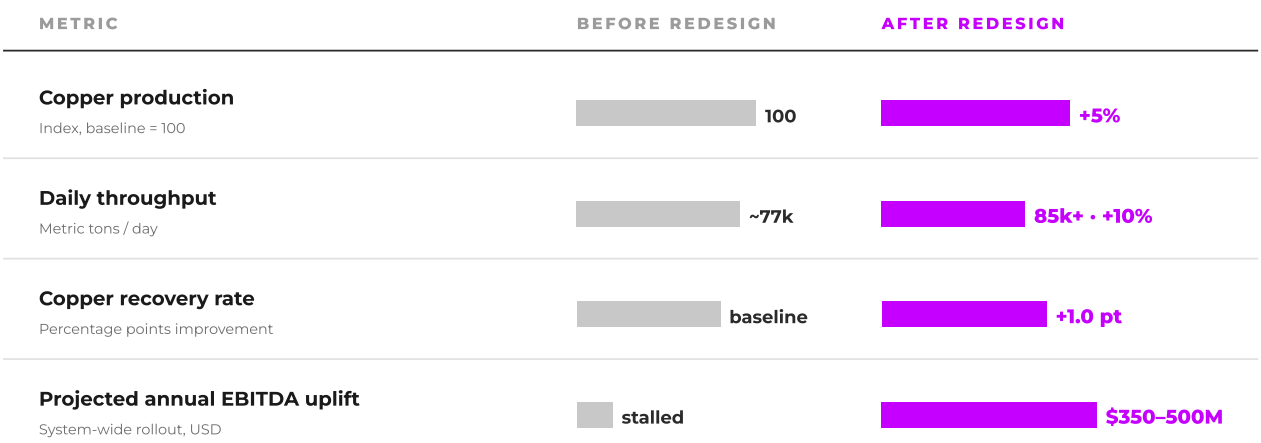

SOURCE · MCKINSEY, "FREEPORT-MCMORAN: UNLOCKING NEW MINING PRODUCTION THROUGH AI TRANSFORMATION"

"Freeport didn't have better AI. They had a better operating model around the AI."

SECTION 02 · TRAIN THE MONKEY FIRST

SECTION

03

Why 94 % keep building pedestals.

Buying servers is fast, tangible, and politically safe. Restructuring a team is none of those things — which is why even executives who can name the right move don’t make it.

Six percent of organizations are doing this differently. The rest are not.

McKinsey's 2025 data is clarifying. Only 6 % of organizations generate 5 % or more of their EBIT from AI. Those high performers are **3.6× more likely** to have redesigned their organization alongside AI deployment. 55 % of them fundamentally reworked workflows when deploying AI. The other 94 % did not. BCG's parallel finding lands in the same place: **60 % of companies generate no material value from AI** despite continued investment.

Figure 03 · Two populations.

The 6% generating >5% of EBIT from AI vs. the 94% generating little or none.

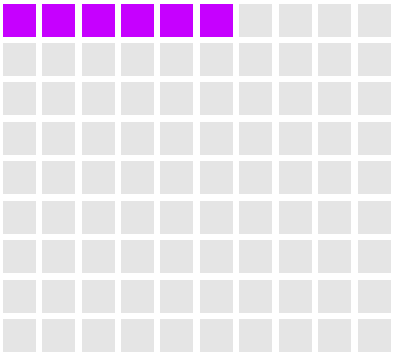

THE 6% · HIGH-PERFORMERS

6%

Generate >5% of EBIT from AI.

3.6× **More likely to have redesigned** the org around AI.

55% **Of high-performers fundamentally** reworked workflows.

60% **Of all companies see no material** AI value (BCG).

SOURCES · MCKINSEY STATE OF AI 2025; BCG "AI AT WORK" 2025

Why the resistance? The political economy of org redesign.

Teller's insight applies directly. Organizational redesign is invisible, slow, uncomfortable, and politically charged. Buying servers is none of those things.

A CTO can present a procurement decision to the board in a single slide. A CHRO explaining why twelve roles need to be restructured, three new team configurations need to be piloted, and decision-making authority needs to shift from function heads to cross-functional leads — that presentation doesn't fit on a slide. It fits in a transformation program that takes quarters, not weeks.

The pedestal is fast, tangible, and easy to approve. The monkey is slow, ambiguous, and requires leadership to make uncomfortable structural calls.

93 % of executives surveyed cite culture and organizational readiness as barriers to AI progress (HBR / PYMNTS, 2025). They know the monkey exists. They fund the pedestal anyway — because that's the line item they know how to manage.

This is not a technology problem. It is an incentive problem. Capital expenditure has a one-line approval. Restructuring twelve roles does not.

Until the cost of *not* redesigning becomes legible — in stalled pilots, in unrealized EBIT, in vendors invoicing for shelf-ware — the pedestal will keep getting funded.

THE DIAGNOSTIC QUESTION

If you stopped buying compute tomorrow, would your AI program still produce a different P&L next year?

If the honest answer is "no — we have nothing to deploy more of," the binding constraint is not infrastructure. It is the operating model that surrounds it. That is the monkey, and it is the only thing standing between current investment and measurable EBIT.

SECTION

04

The six places the monkey hides.

A structural diagnostic across the six domains where AI operating-model redesign is non-negotiable. Each one is harder than buying compute. Each one is required for AI to generate sustained business impact.

Six domains, six monkeys, one diagnostic.

In organizational design for AI, there are six domains where the monkey-first principle applies. Each one requires structural decisions that most AI programs skip. Run your program against the diagnostic question in each card.

01	02	03	04	05	06
Strategy & Value	Structure	Decisions	Workflow	Capabilities	Governance

DOMAIN 01 · STRATEGY & VALUE

Which use cases actually move EBIT?

Most companies have 30–50 AI pilots. Fewer than five are tied to a measurable business outcome with an owner and a deadline. The rest accumulate as a portfolio of demos, which produces vendor invoices and almost no shareholder value.

Diagnostic · Can you name the three AI initiatives that will generate the most value in the next twelve months — and does each have a named executive sponsor?

DOMAIN 02 · STRUCTURE

Where does AI capability live in the org?

In IT? In a center of excellence? Distributed across business units? This is a design decision with massive implications for speed, governance, and talent retention. Most organizations haven't made it deliberately — the structure emerged by accident as different teams started experimenting.

Diagnostic · If asked to draw your org's AI structure on a whiteboard, would the answer match across the CTO, CHRO and the line-of-business leaders?

DOMAIN 03 · DECISION ARCHITECTURE

Who is authorized to act on AI outputs?

This is the domain that separates organizations that scale from organizations that pilot forever. If a supply-chain model recommends reallocating inventory across three warehouses, who acts? In most companies, the answer is: nobody, because no one has the formal authority to trust and execute an AI-generated decision.

Diagnostic · For your top three AI use cases, name the person whose calendar gets a meeting if the AI output is wrong.

DOMAIN 04 · PROCESS & WORKFLOW

Is the AI output in the critical path?

This is the domain McKinsey's data points to most directly. 80 % haven't redesigned workflows. In practice this means AI sits alongside existing processes as an optional input — a dashboard nobody checks, a recommendation nobody acts on, a copilot nobody has been trained to use. Workflow redesign means the process doesn't work without it.

Diagnostic · If your AI system went down for a week, would anyone's daily workflow break? If not, AI is not yet in the critical path.

The two domains that quietly determine whether the rest can scale.

DOMAIN 05 · CAPABILITIES & CULTURE

Can the workforce actually run an AI-augmented model?

You can't run an AI-augmented operating model with the same skills profile that ran the pre-AI version. Data literacy, prompt engineering, human-AI collaboration protocols, algorithmic risk awareness — these aren't nice-to-have training modules. They're prerequisites for the operating model to function.

Diagnostic · The EU AI Act's Article 4 AI-literacy obligation has been in force since February 2025. If audited tomorrow, would your organization pass?

DOMAIN 06 · SYSTEM GOVERNANCE

Who classifies, monitors and owns the AI inventory?

Without system governance, every AI deployment is a liability — legal, reputational, and operational. Governance is the connective tissue that makes all other domains safe to operate. It is not bureaucracy — it is the precondition for scaling without compounding risk.

Diagnostic · Can you produce, today, an inventory of every AI system in production with its owner, risk classification, and monitoring cadence?

A framework for asking the right question first.

The six domains are not a checklist to complete in sequence; they are a diagnostic to apply continuously. The structural test is simple: in any AI investment decision, the operating-model question must be answered *before* the technology question. The order of operations is the operating model.

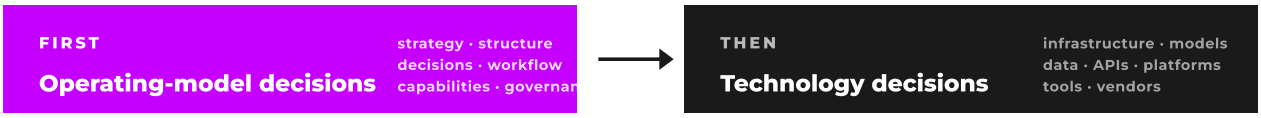

SECTION

05

What Monday morning looks like.

A 90-day sequencing plan for moving from pilots to scaled impact. Four phases, one principle: redesign the operating model around one initiative before scaling to many.

If you recognize your organization in the 94 %, here's the sequencing.

Four phases. The first three weeks produce no new technology. They produce clarity — about what already exists, who owns it, and which workflows are ready to be redesigned around AI as a critical input rather than an optional dashboard.

WEEK 01

01

Map what you actually have.

Inventory every AI initiative currently running. For each one, answer three questions: **(a)** who owns the business outcome? **(b)** Has the workflow been redesigned to make AI a critical input — not an optional dashboard? **(c)** Does someone have formal decision authority based on AI outputs? Most organizations discover the answers are "nobody," "no," and "no." That is the diagnostic.

WEEK 02–04

02

Pick one initiative. Redesign the operating model around it.

Choose the initiative closest to business impact. Assemble a cross-functional team — not a committee, a team with a product manager, shared OKRs, and a sprint cadence. Redesign the workflow so the AI output is in the critical path. Assign decision rights explicitly. This is what Freeport did at Bagdad. It is what the 6 % have in common.

MONTH 02–03

03

Build the governance spine.

Classify your AI systems. Assign ownership for the inventory. Establish a reporting line to the board or executive committee. This isn't bureaucracy — it's the precondition for scaling. Without governance, every additional AI deployment increases risk faster than it increases value.

ONGOING

04

Reallocate budget — invert the ratio.

If your AI budget is 90 % technology and 10 % organizational change, the change-management line is where the next dollar of value comes from. The infrastructure will keep getting cheaper. The operating-model redesign will not get easier by waiting.

The pedestal is already built.

The uncomfortable truth for most enterprise AI programs in 2026: the pedestal is done. The compute exists. The models work. The APIs are available. The infrastructure has never been more accessible or more affordable.

The monkey — the organizational redesign, the workflow restructuring, the decision rights, the role definitions, the governance spine — is the only thing standing between your AI investments and measurable EBIT impact.

Freeport-McMoRan proved it. McKinsey's data confirms it. Gartner's $2.52 trillion price tag quantifies the scale of the misallocation.

Astro Teller's question is the right diagnostic: *what's the monkey in your AI program?* If nobody in the room can answer that question, you're building a very expensive pedestal.

"If nobody in the room can name the monkey in your AI program, you're building a very expensive pedestal."

CLOSING ARGUMENT · TRAIN THE MONKEY FIRST

ENGAGEMENT

If you're mapping AI decision rights, or moving from pilots to scaled impact, reach out directly.

HandsOn works with leadership teams on exactly this transition — from pedestal-building to monkey-training: structuring AI operating models, assigning decision rights, and designing the workflows and governance spine that make sustained EBIT impact possible.

ABOUT THE AUTHOR

HANDSON

Maximilian Stein.

Founder & Managing Partner, HandsOn

Maximilian Stein advises leadership teams on the operating-model redesign required to translate AI investment into measurable EBIT impact. His work focuses on the structural decisions — decision architecture, workflow design, capability building, and system governance — that separate the 6 % of organizations generating real AI value from the 94 % that don't.

HandsOn is an independent advisory firm working with European mid-market and enterprise leadership teams on AI organizational design. The Insight Series publishes original research and field-derived frameworks on the strategy, structure, and governance of AI inside large organizations.

SOURCES

- **Gartner** · Worldwide AI spending forecast, January 2026.
- **McKinsey** · The State of AI 2025 (QuantumBlack).
- **McKinsey** · Freeport-McMoRan: Unlocking new mining production through AI transformation.
- **BCG** · AI at Work: Momentum Builds but Gaps Remain, 2025.
- **EC-Council** · Enterprise AI budget traps, 2025.
- **HBR / PYMNTS** · Enterprise AI readiness gap, 2025.
- **X / Astro Teller** · "Tackle the monkey first."

CONTACT
info@wearehandson.de
wearehandson.de

SERIES
HandsOn Insight
Pillar 02 · Org Design for AI

EDITION
Vol. 01 / № 04
May 2026