


Scaling the System



The “Now Everyone Wants In” Moment

It was the kind of problem Ella never thought she'd have.

Her AI pilot had worked. Not flawlessly. There were hallucinations, odd phrasing, the occasional missed nuance, but it had saved hours and actually made people smile. That alone felt like a win.

What she hadn't expected was what came next.

It started with Procurement. Then Marketing. Then someone from Logistics she'd never met. And then a message appeared directly on Teams:

“Hey Ella, heard AI wrote your last survey summary? Any chance it can help us with our stocktake comms?”

By the end of the week, she had 27 new emails, 14 Teams chats, and 3 calendar invites, all with the same subject line in one form or another:

“Can we use AI for...”

She stared at her inbox. It was flattering. Encouraging, even. But also... a little terrifying.

Because none of these teams had any formal guidance. No consistent tools. No understanding of the risks. And definitely no capacity from Ella's side to support them all.

She opened her notebook, scribbled the word “scale”, and underlined it three times.

Then she added, “without losing our minds.”

That afternoon, she caught up with Leila, the policy lead and her unofficial sounding board.

“We've gone from ‘what's AI?’ to ‘we want in!’ in a matter of weeks,” Ella said. “It's amazing, but it's chaos. Everyone's skipping the boring bits like risk and support.”

Leila nodded. “Because we've made it look easy. Safe. Useful. That's on us.”

“Yeah,” Ella sighed. “And now we either control it and slow things down or let it run wild and hope nothing breaks.”

Leila paused. “What if there’s a third option?”

Ella raised an eyebrow.

“What if we don’t scale the tool first? What if we scale the thinking?”


That night, Ella stayed late. She opened a clean slide deck and wrote the title:

“Scaling AI: What Comes Before the Tools”

She didn’t have the answers yet.

But the question was clear:

How do you scale AI without becoming the department of ‘No’?



Key Learning: Scaling interest isn’t the same as scaling capability. When demand for AI grows, the first priority isn’t more tools, it’s more understanding.

The Governance Gap

By Wednesday, Ella had four teams piloting different AI tools, none of which she'd recommended.

One used a free browser extension to summarise contracts. Another plugged sensitive HR data into an overseas-hosted chatbot. A third asked an AI to rewrite supplier feedback "with a more positive spin," then sent it straight to the vendor. And one enterprising team built a workflow connecting ChatGPT to their inbox using Zapier.

When Ella found out, her stomach dropped.

None of it was malicious. In fact, it was all driven by the right instincts: improve, simplify, save time. But it was also messy, uncontrolled, and borderline dangerous.

She pulled up her old usage inventory. Back then, shadow use was invisible. Now, it was highly visible, and moving faster than governance could keep up.

That afternoon, she booked a meeting with Risk, IT, Legal, and Procurement. The invite said simply:

"AI: We Need a Framework."

At the meeting, she started with humility.

"We don't need red tape. But we do need shared boundaries. Right now, AI is being adopted faster than we can track. That's good and bad."

Ravi from IT chimed in. "We don't even have a clear definition of what counts as AI use."

Legal nodded. "And no consistent approvals process. People don't know when they need us, until something goes wrong."

Procurement raised a hand. "Can we at least agree: no paid tools unless vetted centrally?"

Ella wrote it on the board:


“One: All external AI tools must go through procurement and risk check.”

By the end of the hour, they’d mapped five lightweight guardrails:

1. **Tool Registration:** If you’re using an AI tool, log it.
2. **Data Risk Check:** If it touches sensitive info, get a review.
3. **Human Review:** No AI output should be sent externally without oversight.
4. **Training First:** You must complete AI basics before trialling tools.
5. **Experimentation Space:** A safe, internal sandbox for trying things out.

It wasn’t complex. It wasn’t bureaucratic. It was clear.

Ella smiled. “This isn’t governance to say ‘no.’ It’s governance so we can keep saying ‘yes’, safely.”



Key Learning: Good governance doesn’t mean control. It means clarity and confidence. The faster AI moves, the more essential it becomes to define how, when, and why it should be used.

The Procurement Tangle

Ella was halfway through her morning coffee when her phone buzzed.

“Hey, just a heads up. The Sales team signed up for an AI proposal writer, \$1,200 annual subscription. We only found out when the invoice came through.”

She blinked.

No one had mentioned it in the AI tool registry. No risk review. No IT assessment. And now it was live, processing customer data.

She pulled up the tool’s privacy policy. It stored prompts. It logged user data. It had servers in another jurisdiction. Her stomach turned again.

Ella messaged Leila.

“We’ve officially crossed into ‘procure now, check later’ territory.”

By lunchtime, Ella was in a room with Procurement, Legal, and a red-faced Sales manager who insisted:

“It’s just a tool. It helps with proposals. We didn’t think we needed to escalate.”

Procurement wasn’t angry, just tired. They’d seen this movie before. SaaS tools were easy to buy, harder to unwind.

Legal asked one question:

“What did it do with the last five prompts?”

The Sales manager wasn’t sure. That was the answer they feared most.

They sketched one together on a whiteboard:

- Is it AI-powered?
- Will it handle sensitive or customer data?
- Is it hosted outside Australia?
- Will it be used in decision-making or communications?
- Have you checked the Terms of Use?

If the answer was yes to any of the above, escalation was mandatory.

They called it: AI Procurement Pre-Check.


“Think of it like sunscreen,” Ella said. “You won’t stop the sun, but you can stop the burn.”

They shared the checklist across the org with a clear message:

“If it looks like AI, check with us first. We’ll help you move fast, safely.”

It worked. The next time someone wanted to trial an AI tool, they brought Ella coffee and the checklist.

Progress.



Key Learning: Scaling AI requires more than good ideas. It requires procurement paths that balance speed with safety. A clear intake process prevents reactive firefighting later.

The “We Need a Platform” Debate

The email subject line was polite, but firm:

“Consolidation Proposal – AI Tools Discussion”

Ella opened it. The CIO was asking whether they should standardise around a single AI platform, “preferably Microsoft Copilot, since we already pay for it.”

She wasn’t surprised. The IT team had been quietly tracking the growing sprawl: browser extensions, freemium tools, niche apps with unknown security settings. From a tech risk perspective, the desire to consolidate made perfect sense.

From a business perspective? Less so.

Ella called a cross-functional meeting. IT, business leads, Cyber, and a couple of team-level AI champions. She opened with a map: a visual of 18 different AI tools currently in use across the organisation.

Tom from Procurement laughed. “I didn’t even know some of these existed.”

Leila from Policy raised an eyebrow. “Half of these would have been blocked if they went through proper review.”

But Marketing pushed back.

“The ones that work best for us are small and fast. I’m not convinced a Microsoft tool can do slogan generation like our current app.”

And HR weighed in:

“We just learned how to prompt ChatGPT properly. Don’t make us start again.”

The debate unfolded in real-time. Ella wrote two columns on the whiteboard:

Control vs Capability

IT argued for control. Fewer tools meant easier patching, better visibility, simpler training, and tighter integration.

Business argued for capability. The freedom to choose tools that actually worked for their unique use cases.

Ravi, ever the pragmatist, finally cut through:

“This isn’t about one tool to rule them all. It’s about defining the core stack and knowing when to allow exceptions.”

Ella nodded. That felt right.


They agreed on a principle:

- **Primary AI stack:** Microsoft Copilot and approved in-house tools
- **Requestable exceptions:** Business cases for external tools reviewed quarterly
- **Blocked by default:** Anything high-risk, unvetted, or consumer-grade

Ella summarised it later as:

“Standardise where it makes sense. Permit where it adds value. Block where it creates risk.”

The debate didn't disappear, but now, it had shape.



Key Learning: Standardisation should never kill innovation, but fragmentation will kill trust. A core AI platform creates stability, but exceptions need structured pathways.

The Skills Mismatch

It hit Ella during an “AI Huddle” in the break room.

Tom had just demoed how he used a prompt to rewrite a policy summary. It was clear, snappy, and well-structured, but it left out two compliance clauses. Leila frowned.

“It sounds right... but it’s legally wrong.”

Tom shrugged. “I didn’t change anything. I just asked it to summarise.”

Ella gently stepped in.

“That’s the problem. We’re trusting the output because it reads well, not because it’s correct.”

Heads nodded. Slowly.

Later that day, Ravi messaged her.

“Half our users don’t realise AI tools can hallucinate. Some think it’s just a better Google.”

Ella had been so focused on tool governance, she’d missed something bigger: skill gaps. Not in technology, in literacy.

People knew how to use the tools. But not how to think with them.

She asked HR for help.

They looked at her with sympathy. “We can help run training,” they said. “But what are we teaching? Prompt writing? Data privacy? Model architecture?”

Ella smiled, thinly. “Yes. But also no.”

What she needed wasn’t a workshop. It was everyday fluency, the kind of instinct that helped someone spot when AI got it wrong, even when it sounded right.

She and Leila created a new format: Prompt & Review Clinics.

Short, low-pressure, weekly sessions.

Format:

- One task (e.g. summarise a policy)
- One prompt used
- One human-led review
- One lesson captured

In week one, they reviewed:

Prompt: “Summarise this travel policy into bullet points.”

Output: Accurate, except it skipped the clause about international quarantine rules.

Lesson: “AI drops exceptions. Humans must catch the edge cases.”

They kept it simple. They didn’t blame. They just showed what good looked like.

Soon, people were submitting their own prompts for review.

Ella created a microsite:

“**Better Prompts, Better Thinking**”

Each week, they uploaded a new example. What worked. What failed. Why it mattered.

She realised something important:

The best prompt wasn’t the cleverest.

It was the clearest.

Key Learning: Access to AI tools isn’t the same as capability. True scale requires fluency, not just in using AI, but in reviewing it, questioning it, and knowing where it falls short.

The Metrics Maze

The CFO wanted numbers.

Ella had just finished walking through a few recent AI use cases, time savings, improved summaries, reduced email backlog, when he asked:

“That’s all great. But what’s the ROI?”

She paused.

How do you measure value when the win is someone getting 20 minutes of their day back?

She mumbled something about productivity uplift and avoided making eye contact with the Head of Finance, who was already scribbling into a notebook.

Back at her desk, she opened a spreadsheet. She listed out all the AI pilots so far. Each had anecdotal feedback, a rough time saving estimate, maybe a quote from a user.

But it felt... soft.

She messaged Ravi.

“Any chance we can track AI usage across tools?”

“Technically yes. But does usage = value?”

Ella groaned. That was the whole problem.

She walked down to HR.

“How do we measure things like confidence? Or engagement with new tools?”

HR smiled. “We use proxy indicators. Participation. Sentiment. Peer coaching.”

That gave her an idea.

Ella reframed the problem:

Not everything valuable is quantifiable, but everything valuable leaves a trace.

She created a two-tier metrics dashboard:

Tier 1: Hard Metrics

- Time saved (manual vs. AI output time)
- Accuracy (AI vs. human-reviewed outputs)
- Risk incidents (number, type, mitigation success)
- Tool usage (frequency, breadth across teams)

Tier 2: Soft Signals

- Uplift in innovation survey responses
- Increase in AI Huddle attendance
- Peer coaching sessions initiated
- Voluntary prompt submissions for review

The dashboard didn't promise a dollar figure.


But it showed momentum, depth, and breadth, and it showed improvement over time.

At the next leadership meeting, Ella walked them through both tiers. She ended with a quote from a team lead:

"AI hasn't saved me hours every week. It's saved me mental hours. That matters."

The CFO nodded slowly.

"Not traditional ROI. But maybe it's the right kind for where we are."



Key Learning: You can't manage what you don't measure. But not all value is financial. To scale AI responsibly, organisations need to track impact in both hard metrics and human signals.

The First Real Failure

It was only one sentence. But it made it into the board report.

The monthly operations summary had been generated using an AI draft. The exec assistant, rushing to finalise the pack, skimmed the output and sent it through without a full review.

The problem? The AI had invented a KPI.

“Supplier turnaround improved by 18%.”

Except... it hadn't.

Procurement noticed. They raised it discreetly. But the damage was done, it had gone to the full executive team. No harm operationally, but it shook confidence.

Ella found out an hour later.

She walked straight into the assistant's office. The poor woman looked stricken.

“I'm so sorry. I didn't mean to skip the check, I just trusted it. It sounded right.”

Ella didn't shout. She didn't even sigh. She just nodded.

Because this wasn't just an error. It was a pattern, one she'd been worried about for weeks:

The more human the AI sounded, the less human oversight people applied.

Later that day, she joined a Risk meeting.

Tom was already on the back foot. “Are we going to ban this now?”

Ella shook her head.

“No. But we do need to formalise something we've tiptoed around: a review threshold.”

She proposed a new rule:

Any AI-generated content used in reporting, policy, or decision-making must go through named human review.

They gave it a name: “**AI Output Assurance**”.

A real person. A real name. A real check.

Not optional. Not assumed. Not a checkbox.

Then she updated the AI in Practice guide with a new section: **First Real Failure – What We Learned**

She included:

- The prompt used
- The output
- The incorrect sentence
- The checklist that should have caught it
- What they’d changed since

No blame. Just facts.

The assistant even offered a quote:

“It’s not AI’s fault. It’s mine for skipping the step. I won’t make that mistake again.”

That quote did more than the policy ever could.



Key Learning: Trust in AI must be earned and verified. When outputs sound confident, humans must become more vigilant, not less. Failure isn’t the end; it’s the foundation for stronger controls.

The Operating Rhythm

Ella sat staring at her calendar.

It was full. Not with meetings, but with chaos.

Everyone wanted her time. One team wanted help vetting an AI tool. Another was piloting something and needed a “quick review.” HR was planning a training session. Legal had questions about LLM licensing.

She had become the bottleneck.

What started as a movement had turned into a maze, and she was the only one holding the map.

That’s when she realised: AI wasn’t the problem. It was the rhythm. Or lack of one.

She pulled out a blank sheet and wrote one word at the top: Cadence.

Ella didn’t want to centralise everything. But she needed structure. A beat. Something predictable.

She sketched out a new operating rhythm:

Weekly

- **AI Huddles:** informal, open sessions for sharing learnings and blockers
- **Tool Triage:** quick 20-min review of new requests, logged for risk

Fortnightly

- **Use Case Clinic:** deep dives into active pilots with feedback and peer review

Monthly

- **Governance Touchpoint:** cross-functional check-in with IT, Risk, Legal
- **AI Dashboard Update:** circulate key metrics and soft signals

Quarterly

- **Working Group Summit:** showcase of wins, failures, and emerging themes

She sent the proposal to her leadership sponsor.

The reply came back an hour later:

“Approved. Looks clean, scalable, and not too heavy.”

That Friday, during the first “AI Huddle” under the new rhythm, something happened.

People stopped asking for help.


They started offering it.

Someone shared a great prompt template. Someone else offered to run the next session. A quiet analyst shared an unexpected AI success in file classification.

The shift was subtle, but profound.

They didn’t just have tools.

They had momentum and a way to maintain it.



Key Learning: AI adoption at scale needs more than tools and policies. It needs rhythm. A lightweight, repeatable operating model turns chaos into continuity.

The Decentralisation Question

Ella stared at her to-do list. It had tasks from seven different teams. Four pilots. Two tool reviews. One legal concern. And a request to run a training session for Finance.

She rubbed her eyes.

“This isn’t scaling,” she muttered. “This is stretching.”

Later that day, during the fortnightly Use Case Clinic, she floated an idea.

“What if... I wasn’t the only one doing this?”

The silence was immediate. Then Ravi spoke.

“You mean like... more Ella’s?”

She grinned. “More like Ella-ish.”

She pitched the concept at the next Working Group Summit:

“We need to decentralise. Not the risk. Not the governance. But the enablement.”

She proposed the creation of AI Champions, volunteers from each major function trained to:

- Support prompt design and tool usage
- Monitor local experimentation
- Liaise with the central AI working group
- Share lessons from the front line

Not owners. Not police. Just enablers with context.

The idea landed better than expected. People were already informally playing that role, they just hadn’t been recognised.

Within two weeks, she had nominees from every function.

She built a short induction pack:

- AI principles
- Approved tools
- Risk triggers
- Prompt best practices
- Who to contact when things go wrong

She scheduled monthly AI Champion check-ins and added them to the AI Huddle calendar as co-hosts.

The first time a champion corrected a prompt before it went live, Ella got the message.

“AI isn’t centralised anymore. It’s distributed, responsibly.”

And that was the point.

Ella’s job was no longer to answer every question.

It was to make sure the right people knew how to answer them.

Key Learning: True scale happens when capability is distributed. AI Champions bridge the gap between central governance and everyday use, embedding responsibility without creating red tape.

The Scalable Success

The HR team had a problem.

Every new hire received a 12-page onboarding document filled with corporate jargon, compliance clauses, and acronyms no one explained. Most skimmed it. Some ignored it. A few admitted they didn't read it at all.

Ella listened quietly in their check-in.

Then Leila from HR spoke up.

"Could we... rewrite it with AI? Make it more human?"

There were nods, cautious, but hopeful.

Ella didn't promise a magic fix. But she offered to co-design the pilot.

They took the original onboarding doc and broke it into sections. For each, they drafted prompts like:

"Rewrite the following in plain language suitable for someone on their first day of work."

They included rules:

- **No hallucinating**
- **No skipping compliance clauses**
- **Always flag changed wording for review**

Each section was reviewed by HR, Legal, and a new hire for clarity.

Where the AI oversimplified, they adjusted the prompt.

Where it omitted something important, they added a red flag.

They called it the AI Rewrite Cycle:

1. Draft with prompt
2. Human review
3. Feedback loop
4. Approved variation
5. Logged prompt + checklist

The final result?

A five-page, readable, friendly onboarding doc, with the same integrity, minus the corporate fog.

They embedded a note at the top:

“This guide was co-written by AI and reviewed by real humans. Because clarity matters.”

HR was thrilled. New hires responded positively. Legal gave a quiet thumbs-up.

Ella documented the whole thing:

“Scalable Success: AI-Assisted Onboarding Rewrite”

She included:

- Prompts used
- Review templates
- Risk flags encountered
- Time savings
- Feedback from end users

It wasn't flashy. But it was repeatable.

“Now this,” she told the AI Working Group, “is something worth scaling.”



Key Learning: Scalable success doesn't come from novelty. It comes from structured, human-reviewed wins that others can adopt with confidence.

Repeatability is the real multiplier.

The Leadership Workshop

The room was unusually quiet.

The Executive Team had gathered for a half-day workshop on AI. Ella stood at the front, flanked by a whiteboard and three post-it-covered boards. There was no glossy deck. No hype video. Just stories, metrics, and tension.

The COO leaned forward.

“Look, I get the pilots. But when are we going to move faster?”

The General Counsel cut in.

“And when are we going to put the brakes on some of these experiments? We’re walking into a regulatory storm.”

Ella nodded.

She had expected this.

“This is the real conversation,” she said. “It’s not about pilots anymore. It’s about who we want to be.”

She drew two axes on the whiteboard.

- **Y-axis: Speed**
- **X-axis: Control**

She plotted the current state: fast experimentation, loose coordination.

She plotted where they were headed: structured enablement, embedded risk controls.

Then she posed the question:

“What are you solving for? Cost? Efficiency? Differentiation? Reputation?”

The silence that followed wasn’t awkward, it was thoughtful.

Tom broke it.

“We need to be bold. If we’re not fast, we fall behind.”

Leila followed.

“We also need to be trusted. If we lose public confidence, no amount of speed matters.”

Ella added a new quadrant:

“Safe Acceleration”, high speed, high control.

“That’s what we’re aiming for. Not safety or scale. Both.”

She walked them through the foundations already in place:

- Use case clinics
- Prompt review cycles
- AI Champions
- Approved tool registry
- Real-time feedback from pilots
- Risk oversight embedded from day one

Then she showed the soft signals:

- 12% increase in innovation scores
- 6 teams initiating their own AI onboarding
- First repeatable, cross-functional use case in HR

By the end of the session, no one asked, “Should we do AI?”

They asked, “How do we do it with integrity and speed?”

Ella left with three executive asks:

1. **Formalise the AI Working Group as a permanent fixture**
2. **Allocate baseline funding for tools, training, and ops**
3. **Build an internal comms plan to reinforce principles, not just tech**

She walked out into the hallway, exhaled slowly, and smiled.

They weren’t just onboard.

They were aligned.



Key Learning: At scale, AI stops being a tech project and becomes a strategic decision. Real alignment comes not from consensus, but from shared clarity about risk, value, and direction.

The Second Inventory

Three months later, Ella opened the same spreadsheet she had started a year ago.

Back then, it had been a messy list: tools people were experimenting with, use cases she'd overheard, and rough estimates scribbled in the margins. Shadow AI, untracked prompts, blind risk.

Now, she clicked on a new tab:

“AI Use Inventory – Reviewed & Active”

It was clean. Structured. Governed.

Each row told a story:

- The department using it
- The tool's name
- Purpose and benefit
- Review date
- Champion assigned
- Risk class
- Review outcome

But it wasn't the neatness that struck her.

It was the shift in tone.

People weren't hiding their usage anymore. They were proud of it, and accountable for it.

She scrolled to the final column: “Repeatable?”

Nearly half the use cases were marked Yes.

She clicked into the comments:

- “Successfully scaled to 3 teams”
- “Prompt template shared org-wide”
- “Now part of new hire induction”

It wasn't perfect. But it was progress.

Leila stopped by her desk.

“You realise you’ve built something rare, right? Most companies flail at this point.”

Ella shrugged. “I didn’t build it alone.”

“Still. Most people try to scale tools. You scaled thinking.”

That stuck with her.

Later that afternoon, she ran the final AI Huddle for the quarter. The room was full. Not with curiosity anymore, but with capability.

The AI Champions ran the session. Ella didn’t even have to present.

She just watched.

One prompt.


One review.

One discussion.

One lesson.

Exactly how it started.

But now, it was theirs.



Key Learning: Scaling AI isn’t just about reaching more people. It’s about reaching maturity. When the thinking scales, the tools follow. And the organisation learns to lead itself.

Credits & Copyright

Author

Dharshun Sridharan, Aresia

Illustrations / Image Creation

Generated with assistance from ChatGPT (OpenAI)

Contact

dsridharan@aresia.com

© 2025 Aresia. All rights reserved. No part of this work may be reproduced, distributed, or transmitted in any form or by any means without prior written permission.