



BURAK GALİBA
Thoughts on AI Transformation

ESSAY · ENTERPRISE AI

Applause Is Not an Operating Model — How AI Demos Get Declared "Ready" Before the Real Work Starts

The demo was flawless because the data was frozen. How AI demos get declared "ready" at exec meetings — and the intake process that would stop it.

AI STRATEGY

LEADERSHIP

TRANSFORMATION

Burak Galiba · August 2, 2026 · 6 min read

<https://burakgaliba.com/blog/applause-is-not-an-operating-model>



SCAN TO READ ONLINE

I've watched this film enough times, in enough different enterprises, that I can narrate the next scene before it plays. The company changes. The plot never does.

It opens with a directive from the top: **we need AI**. Visible progress, this quarter. No definition of what, no process for how — just pressure, distributed evenly across the org chart. Every business unit hears the same instruction: show something.

And nothing stops them, because nothing exists to stop them. The AI licenses are already paid for; the API tokens are already provisioned. A demo needs no budget request, no procurement gate, no review of any kind — it is free and permissionless. So demos bloom in every corner of the building at once.

Picture one of those teams — a composite, but a faithful one. Say it's the group that answers policy questions for claims handlers. They export the procedure manuals to PDF, pull three systems into spreadsheets, and wire an agent on top. Six weeks later it runs at an executive meeting, and it is flawless — every question answered instantly, correctly, with a citation. Applause. The team gets named at the next town hall. By Friday, the assistant is being advertised internally as ready.

That word — *ready* — is where the damage starts.

Perfection is a property of frozen data

Here's what the applause was actually applauding: a system answering questions about a hand-curated snapshot of reality.

The demo is flawless *because* the data is frozen, not despite it. A static export has no access-control problem — someone with the right permissions did the exporting. No freshness problem — nothing changes. No conflicting-sources problem — the curator already picked the winners. No edge cases — those were groomed out on the way to the spreadsheet.

Every hard problem of production wasn't solved by that demo. It was deleted by the export button.

None of this is dishonest — prototypes are supposed to fake the hard parts; that's what makes them cheap. The failure comes next, when nobody in the room asks "what data is this running on?" and demo perfection gets read as product maturity.

"Make it live" is a euphemism for "build it"

After the applause, the assistant lands on IT with a simple-sounding request: make it live.

Which turns out to mean: integrate the live systems, build the pipelines that keep data current, implement the permission model, decide what happens when sources disagree, add evals, logging, rollback. In other words — build the product, from scratch. The demo was not version 0.9 of the system; it was a storyboard of it. Almost nothing survives the handoff except the prompt and the idea.

I've written about the [deployment side of this handoff](#) — the paved road that gets business-unit agents into production safely. This essay is about the other end of the pipe: what enters it, and who decided it should.

Meanwhile, the snapshot rots

While engineering scopes the rebuild, the demo version is still in users' hands — it was advertised as ready, so people use it.

And the snapshot ages from the moment of export. Procedures get revised; the PDFs don't. Records change; the spreadsheets don't. The assistant keeps answering from last quarter's reality with the same fluent confidence it had on demo day — the wrong answers arrive dressed identically to the right ones.

An agent confidently serving stale guidance is worse than no agent at all. Users catch one bad answer, start double-checking every answer, then skip the middleman entirely. Adoption decays in weeks — and trust burned this way is the [most expensive thing to rebuild](#), because the next tool inherits the skepticism this one earned.

"It was ready. Why is IT taking so long?"

Now the killer scene. Months into the rebuild, leadership — still anchored to the applause moment — asks the question this whole setup was built to produce: *we saw it working in March. Why is IT taking so long?*

Engineering is being measured against a mirage it never created. And notice that nobody behaved badly. The business team didn't lie; they built exactly what was asked, fast. Leadership applauded a genuinely good idea. IT isn't slow — months of integrations, pipelines, permissions, and evals is normal speed for the real thing. Everyone acted reasonably, and the organization still manufactured a fiasco.

That combination — reasonable people, terrible outcome — is the signature of a missing operating model.

Applause is the intake process

Look at what structurally happened. There was no intake for AI work, so the executive showcase became the intake. There were no prioritization criteria, so applause became the criterion. There was no definition of "ready," so the word was free for anyone to claim.

Recognition rewarded the demo. Funding followed the showcase. Nothing rewarded shipping. Every incentive in the building pointed at producing more demos — which is exactly how [pilot purgatory](#) gets stocked, one standing ovation at a time. A report from MIT's NANDA project found that [95% of enterprise GenAI pilots produce no measurable P&L impact](#); I suspect a good share of that number is applause doing a job an operating model should be doing.

This is not a technology failure — the agent worked. It is not the business unit's fault — they answered the incentives with energy most transformation programs would kill for. It is an operating model failure, and it belongs to whoever said "we need AI" without building the machinery for what happens when someone responds.

Build a front door

The fix is not a bigger committee. It's a small amount of process, placed exactly where the energy enters the system:

- **One front door.** A one-page intake that every AI initiative files *before* it reaches an executive stage: named owner, data sources (static snapshot or live connection?), target workflow, and a value hypothesis [you could someday measure](#). Not a permission slip — a registration.
- **Criteria, not decibels.** Published prioritization criteria and funding gates, so what advances is decided by written rules, not by whoever demos loudest.
- **A published definition of "ready."** A demo on static data is a proposal, not a product. Proposals can earn a production track and engineering capacity. They cannot earn the word *ready* — that's reserved for systems on live data with an owner, evals, and a rollback plan.
- **The showcase rule.** Every demo shown to leadership states, on the slide, what data it runs on and what production would take. Then leadership can do both right things at once: applaud the idea, and fund the rebuild.
- **Redesign the applause.** The team that ships to production gets at least the recognition the demo got. If the town-hall mention only ever goes to demo day, your incentive system is formally instructing people to stop at demo day.

Note what's missing: any punishment for the business units. Their initiative is the scarcest raw material in most AI programs. The goal is never to dam the river — it's to give it a riverbed.

Where to start

Before the next executive showcase, write two artifacts. The one-page intake form — owner, data sources, workflow, value hypothesis. And the showcase rule, one sentence in the meeting invite: every AI demo states its data source and its path to production on the final slide. Neither needs a committee. Both change what the applause means.

If you want to know whether your organization has the machinery or just the enthusiasm, my free [AI Readiness Score](#) takes about ten minutes — 20 questions across pilots, data, talent, and governance, including the operating-model ones: who owns intake, what "ready" means, how funding gets decided. It will tell you whether you've built a front door, or just a stage.

Read the essay online at <https://burakgaliba.com/blog/applause-is-not-an-operating-model>

I built a free AI Readiness Score — 20 questions across pilots, data, talent and governance, about ten minutes, no email required: <https://burakgaliba.com/readiness>

© 2026 Burak Galiba · <https://burakgaliba.com> · Views are my own.