BURAK GALİBA
Thoughts on AI Transformation

ESSAY · ENTERPRISE AI

The Hackathon-to-Production Gap — A Paved Road for the AI Agents Your Business Units Already Built

Business units built AI agents at a hackathon. Now they're stalled at IT. Why the fix is a paved road: tiered access, published SLAs, and real dates.

AGENTIC AI | GOVERNANCE | LEADERSHIP

Burak Galiba · July 26, 2026 · 5 min read
https://burakgaliba.com/blog/hackathon-agents-path-to-production

SCAN TO READ ONLINE

Your business units went to a vendor hackathon and came back with working AI agents. Not slideware — agents that read tickets, draft replies, touch real workflows. The demos landed. The sponsors are excited. And now those agents sit at IT's door, waiting for something that doesn't exist: a way in.

There is no deployment standard for agents. No decided policy on whether they get access to live systems, or at what privilege level. No security review path sized for software that acts. The governance committee has discussed it twice and decided nothing. So the business escalates through side channels and demands hard dates — and IT feels ambushed, asked to grant production access to code it didn't build, on infrastructure it has to defend.

Both sides are right. That's exactly what makes this one dangerous.

This is pilot purgatory with the safety off

You know pilot purgatory: demos that impress and never ship. The hackathon version is the same disease with two complications.

First, the politics. In classic purgatory, a central team built the pilot and quietly shelved it. This time the business built the thing and watched it work. When it stalls at IT, that doesn't read as prudence — it reads as obstruction, and every week of silence converts a sponsor into an escalation.

Second, the artifact. A stalled dashboard is a sunk cost. A stalled agent is a loaded one. Software with tool access can act — write to systems, send messages, move data — so waving it through on goodwill is reckless in a way approving a report never was. IT's caution isn't theater. It's the right instinct pointed at the wrong target.

Stop debating agents. Publish a road.

The way out is not a faster yes or a safer no. It's a paved road: a published, tiered path from prototype to production that any agent — any hackathon, any business unit, any vendor stack — can walk without a bespoke negotiation.

That shifts IT from gatekeeper to platform. A gatekeeper debates each agent on its merits, forever. A platform publishes a standard, a procedure, and dates, then holds everyone — including itself — to them. One standard decision beats fifty case-by-case debates, and it's the only version of control that scales past the third agent.

Four tiers, earned one at a time

The core of the standard is a tiered access model. A shape I recommend:

- **Tier 0 — sandbox.** Read-only, synthetic or masked data. Every hackathon agent lands here by default, immediately, no review required.
- **Tier 1 — read-only production.** Live data, zero writes. Entry requires a named owner, a scope statement, and documented eval results.

- **Tier 2 — write with approval.** The agent proposes; a human approves each action before it executes. Most agents should live here for months.
- **Tier 3 — scoped autonomous write.** Narrow, bounded actions without per-action approval — earned by measured error rates at Tier 2.

Every tier runs on least-privilege service accounts — the agent gets its own identity, never a human's credentials — with audit logs from day one. Graduation criteria are published, so "when do we get more access?" has an answer that isn't a meeting.

One warning: the tiers describe what agents *may* reach. Whether there's anything to reach — most enterprise systems have no API or MCP server for an agent to connect to — is its own problem, and I've written up how to build that missing access layer.

A procedure with a date on it

The other half of the road is a deployment procedure with an SLA.

The business submits a short package: owner, scope, systems and data touched, eval results against the target tier. IT checks it against the published standard — identity, logging, data boundaries, rollback. A decision arrives within a fixed number of business days. Ten is defensible. "The committee will discuss it" is not, because a committee that discusses without dates is a place prototypes go to die.

That reframes the governance committee's job. It is not to review agents one by one — it will never keep up, and Gartner already expects over 40% of agentic AI projects to be canceled by 2027. Its job is to publish the standard, set the SLA, and revise both as evidence accumulates — because AI governance is a deadline, not a discussion format. Anchor the standard on NIST's AI Risk Management Framework: risk tiers, human-oversight gates, and evaluation before promotion, without inventing doctrine from scratch.

Saying nothing doesn't remove the risk

Here's the uncomfortable part: a stalled agent is not a contained agent.

When the official path is blocked or undated, motivated teams route around it. The agent reappears on a personal subscription or a vendor sandbox, pointed at exported spreadsheets — unlogged, unreviewed, invisible. Saying no, or saying nothing, doesn't remove the risk. It relocates it to where you can't see it.

An agent on the paved road — least privilege, human gates, full audit trail — is *safer* than the status quo you're defending. That's the argument to bring to your own security team, and it has the advantage of being true.

This week

Draft the tier table. One page: four tiers, entry criteria for each, and the number of business days a decision takes. You don't need the committee's blessing to write it — you need it as the concrete proposal that turns the next escalation meeting into a ratification.

Then count the agents already waiting at your door. That's your backlog, and every one has a sponsor watching what you do next.

To see whether the rest of the organization can support what the business just built, my free AI Readiness Score takes about ten minutes — 20 questions across pilots, data, talent, and governance — and shows where your road needs paving first. The hackathon proved your people can build agents. The score tells you whether your enterprise can run them.

Read the essay online at https://burakgaliba.com/blog/hackathon-agents-path-to-production

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