tray.ai

FIELD GUIDE

# Five signs you've outgrown n8n

What to do when the thing you built has to become a system the business runs on.

*For the builder who made it work on n8n, and now has to get it past review, into production, and running at scale.*

TRAY.AI

# Table of Contents


The wall between a working pilot and a production system 3

Five signs you've outgrown n8n 4

1. Security won't sign off, and security is now your job
2. Governance stopped being optional, and it's behind the paywall
3. AI agents are multiplying faster than anyone can govern them
4. It strains at scale, and loses records under load
5. When it breaks at midnight, the answer is a forum thread

A production readiness checklist 9

What waiting costs 10

This is the moment teams move to Tray 11

# The wall between a working pilot and a production system

n8n is a good way to build. You can stand up a working automation in an afternoon, wire in an AI agent, and prove a real result before lunch. That's exactly why so many teams start there.

The trouble starts later, when the thing you built has to become a system the business depends on. That's a different job. It's the distance between *it works on my instance* and *it runs in production, under review, at scale, with someone accountable when it breaks*.

If you built it, you're usually the first to feel the ceiling. Here are five signs you've reached it, and what to do about each one.

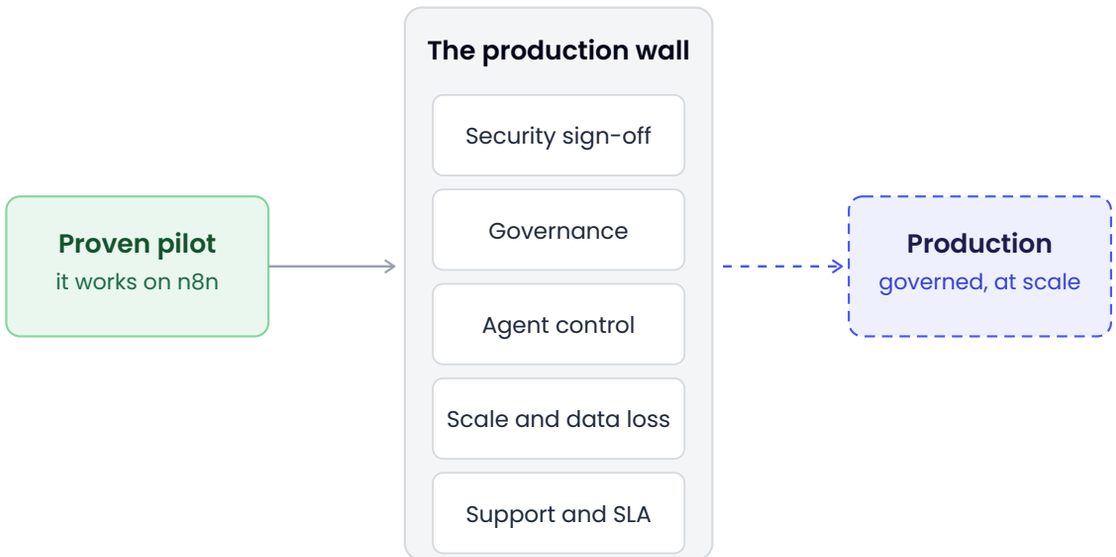


*The pilot works. Five things stand between it and production.*

1

SIGN #1

# Security won't sign off, and security is now your job

Security can be a real pilot killer. Your workflow is solid but nobody will approve a self-hosted instance for production.

It's a legitimate concern. Late 2025 through 2026 brought a run of critical n8n vulnerabilities, including unauthenticated remote code execution, across tens of thousands of internet-exposed instances. One is on CISA's Known Exploited Vulnerabilities catalog, confirmed exploited in the wild.

Self-hosting is what makes patching those vulnerabilities your job: you own the patching, the exposure, and the incident response 24/7.

For a side project, that's manageable. For a system the company relies on, it's a standing liability, and it's usually the reason a security or compliance reviewer says "not like this."

**WHAT TO DO**

When the value is proven but the deployment can't clear review, move the workflows somewhere the patching burden isn't yours and the certifications your reviewer asks for already exist. That's faster than hardening a side project into a platform your team has to defend forever.

### What self-hosting puts on your team

- Patching and CVEs
- Public exposure
- Incident response, 24/7
- Backups and monitoring
- Scaling and uptime
- SOC 1 and SOC 2 evidence

*Self-hosting is a shared-responsibility model, and most of the responsibility is yours.*

*Governance becomes a requirement and your cost jumps in the same moment.*

THE PAYWALL PROBLEM

2 SIGN #2

# Governance stopped being optional, and it's behind the paywall

Early on, "who can edit what" doesn't matter much. Then an auditor, a customer security questionnaire, or your own CISO asks for role-based access, SSO, and a clean audit trail you can stream to your SIEM.

On n8n, those controls sit in the paid tiers, and the 2025/2026 pricing change moved several of them up into a higher Business tier. Teams tell the same story: the moment they needed a SOC 2 report or proper audit controls, they found them gated behind the enterprise plan at a step up in price they hadn't budgeted for. If you run outside your company's core engineering group, it's worse, because proving all of that is on you.

So the moment governance becomes a requirement is also the moment your cost and your operational load both jump.

**WHAT TO DO**

Treat governance as a platform requirement, not a line item you buy later. If RBAC, audit, and log streaming aren't standard, you'll re-implement and re-prove them every time the bar rises.

*Any builder can point an agent at production, and no one can see the calls.*

SHADOW AI, IN PRODUCTION

3 SIGN #3

# AI agents are multiplying faster than anyone can govern them

n8n's AI and MCP features are genuinely strong, and at scale that's a problem. Any builder can stand up an MCP server and point an agent at an internal system. There's no central place to say who can call what, with what scope, and to audit every call after the fact.

It often starts as one person's project. That's fine with three workflows and one builder. It's a liability when the automation everyone now depends on lives in one person's head, connects to production systems, and the person who has to answer for it, your CISO, can't see or control the calls.

**WHAT TO DO**

As agents move into production, put a control layer in front of them: policy, scope, versioning, and a usage audit on every agent and MCP call. Governed agentic execution is the part most teams find they need only after they've shipped a few agents without it.

4

SIGN #4

# It strains at scale, and loses records under load

The patterns that work in a demo start to crack under real volume. Branches that converge unevenly can stall or drop items at the merge step. Execution logs balloon the database until things slow and you're paying for emergency storage. Concurrency becomes a ceiling, and keeping up means load balancers, queue mode, extra workers, Redis, and often your own runners bolted on around n8n. That's a second system to operate on top of the first.

The worst version is silent: records that don't make it through, with no error to tell you. Two branches carry different numbers of items into a merge. Combined by position, the merge pairs item one with item one, two with two, and stops when the shorter branch runs out. The extra items on the longer branch simply fall off.

In a pilot you'd never notice. In production, that's a data-integrity problem with your name on it.

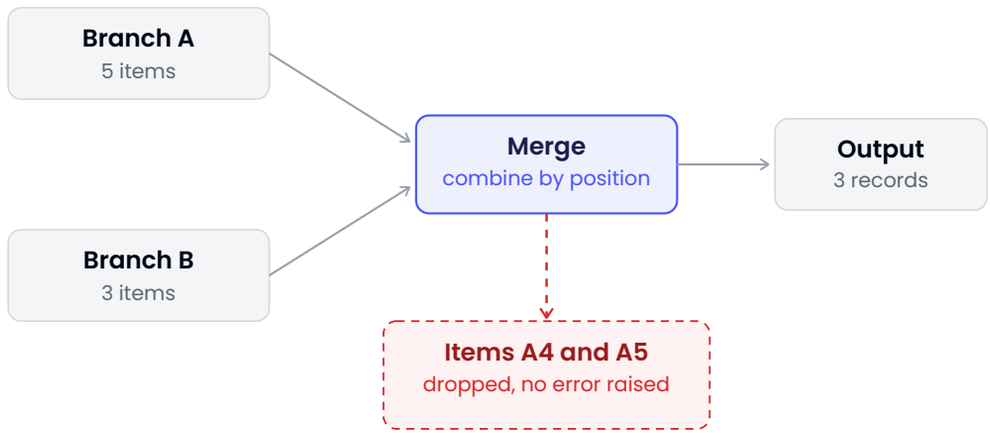


*How an uneven-branch merge loses records: the shorter branch decides where it stops.*

WHAT TO DO

If your workflows join, reshape, and move real volume, you want a managed runtime that scales for you, with streaming and real observability, so large payloads and high volume don't silently drop records, not hand-rolled fixes and a pile of infrastructure you now maintain.

*"Who do we call, and what have they promised us?"*

THE QUESTION PROCUREMENT ALWAYS ASKS

5 SIGN #5

# When it breaks at midnight, the answer is buried in a forum thread

Production systems fail. The question is what happens next. On n8n, guaranteed SLAs and dedicated support sit at the enterprise tier. Below that, your incident plan is a community forum and your own team.

For a system the business runs on, that's the gap that turns a 20-minute outage into a half-day one. "Who do we call, and what have they promised us?" is a question procurement will ask, and "the community" is rarely the answer that clears the deal.

**WHAT TO DO**

If uptime matters, the platform needs a contractual SLA and real support behind it as standard, not a tier you reach only after something has already gone wrong.

SCORE YOURSELF

# A production readiness checklist

Tick every box you can't confidently clear today. Each one is a risk you own the moment this system carries real load, whether you stay on n8n or move.

## SECURITY AND PATCHING

- [ ] When a critical CVE drops, we're not sure who patches our instance, or how fast.
- [ ] We spend real hours each week patching, monitoring, and backing it up.

## RELIABILITY AND SCALE

- [ ] We've seen records drop, hit a concurrency ceiling, or watched the database balloon under load.
- [ ] We can't prove no records are lost when branches merge or volume spikes.

## GOVERNANCE AND CONTROL

- [ ] RBAC, SSO, and audit-log streaming aren't standard for us, or sit behind a higher tier.
- [ ] We couldn't hand an auditor a clean access and audit trail today.
- [ ] Any builder can point an agent at a production system with no central policy, scope, or audit.

## COST AND SUPPORT

- [ ] The real operating cost — infrastructure, ops, and security time — is higher than the license suggests and hard to predict.
- [ ] We don't have a contractual SLA we could point a nervous executive to.

**Three or more ticked?** You've outgrown the build-it-yourself stage. None of this means n8n was the wrong place to start. It means you've proven the work, and the work now needs a production home.

# What waiting costs

The cost of moving will never be lower than it is right now, while the build context is fresh and the footprint is still small. Every month you wait, more workflows go live, more teams build on top, and the migration you could do this quarter becomes the one nobody wants to touch next year.

Waiting is paid in ownership and risk. Every week the platform stays yours to patch, monitor, back up, and answer for. Every new workflow is another dependency to move later. And the longer you wait, the more likely the decision gets made for you, by an audit, a security review, or the next CVE, on a timeline you don't control.

Set against that, the move is small. Teams that handle this well come up for air before a renewal or an audit forces the decision, and they look at it while it's still a week of work rather than a year of it.

> Moving isn't the scary part. **Waiting until you have no choice is.**

# This is the moment teams move to Tray

Tray is the managed, governed platform you run automations and AI agents on in production. And you don't start over: the n8n to Tray Migration Kit reads your n8n exports and moves your workflows onto Tray as native Tray workflows, in minutes, so the thing you already proved can go to production and scale.

### Managed and governed by default

RBAC, audit, and log streaming are standard, not a higher tier and not yours to run.

### Governed agents, built in

The platform sets policy on, and audits, every agent and MCP call. IT can finally say yes to agents in production.

### Security and compliance, handled

SOC 1 and SOC 2 Type 2, HIPAA, GDPR/CCPA, regional data residency. The next CVE isn't your team's fire drill.

### Analyst-recognized

Named a Visionary in the Gartner Magic Quadrant for iPaaS, three years running. n8n isn't evaluated in it.

### Built to scale

A managed runtime that scales for you, with streaming and real observability, so volume and complexity don't mean re-architecting.

### Real support

A contractual SLA and enterprise support, standard, so a midnight incident has a number to call.

**See the gaps before you go to production**

Score your n8n setup on security, compliance, agent governance, reliability, scale, and support.

Get the free checklist →

tray.ai

# Move it to production.

The managed, governed platform you run automations and AI agents on in production. Bring your n8n workflows across in minutes with the n8n to Tray Migration Kit.

Get the free n8n readiness checklist →

A FIELD GUIDE FROM

**Tray.ai**

The leading enterprise orchestration platform for data and AI

Five signs you've outgrown n8n

VERSION 1.0 · JULY 2026

Copyright © Tray.ai Inc. All rights reserved.