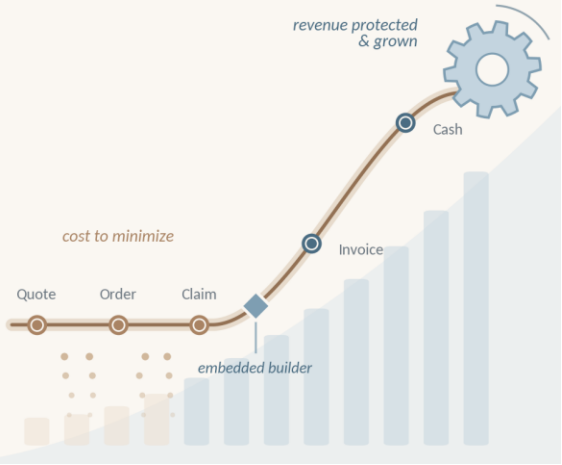
DRAUP RESEARCH BRIEF

OPERATIONS TRANSFORMATION SERIES · SEPTEMBER 2026

From Cost Center to Revenue Engine

How Operations Teams Protect and Grow Revenue in the Age of Agentic AI



Abstract

Operations is usually managed as a cost to be minimized, yet it sits on the revenue path: every quote, order, claim, invoice, and support case either protects revenue or leaks it. Across industries, operations organizations are running more AI pilots than ever while shipping few into production. This brief argues the gap is rarely about technology: it is the distance between the people who understand the work and the people empowered to change it. It examines, through an illustrative composite case, how an enterprise operations organization closed that distance — by starting with honest process discovery and embedding builders directly inside its process teams — and maps the approach onto thirteen industries — from high tech, software, retail, and consumer goods to telecom, biopharma, healthcare, and the public sector.

KEY FINDINGS

- **The constraint is deployment context, not AI capability.** Pilots fail because agents are built to documented processes that no longer match how work actually flows.
- **Embedding closes the distance.** In the case examined, engineers stationed inside operational teams moved 11 of 14 builds into production in year one, versus a historical conversion rate below 30%.
- **Discovery comes first.** Mapping the real process — minute steps and workarounds included — proved the most valued output, even where no automation followed.
- **Pre-agreed governance accelerates.** Deterministic guardrails and graduated autonomy, negotiated once, removed per-project risk negotiation.
- **Operations is a revenue lever, not just a cost line.** The highest-value builds recovered leakage, sped up quotes and approvals, and reduced disputes that delay cash.
- **The model is industry-neutral.** Regulation changes the autonomy ceiling, not the mechanism.
- **The FDE is a talent decision, too.** The scarce skill is translation, and the roles around each build change, so HR co-owns the model.

Prepared by **Vijay Swaminathan**, Co-founder & CEO, Draup

· Who this is for: operations, transformation, and technology leaders in any operations-intensive industry, and the CHROs who partner with them

About the case: *The organization described here is an illustrative composite, not a single named company. Program details and outcome figures are representative and aligned with published industry benchmarks for each process area. See About This Brief.*

> **For CHROs sharing this with your COO:** this brief is written for operations leaders, but the model it describes succeeds or fails on a talent decision — a new, scarce role embedded in teams whose work is about to change. Section 08 sets out that side of the program: the role profile, where to find it, and what happens to the roles around it. It is the conversation to have with your COO before the first engineer is hired.

01 Executive Summary

In the illustrative case this brief examines, an operations organization at the core of a multibillion-dollar enterprise entered 2025 in a familiar position: heavy automation investment, dozens of AI pilots, and almost nothing in production. More than sixty improvement initiatives were in flight; fewer than a third had shipped anything a frontline team used. The organization was not short of technology, budget, or intent. It was short of a *mechanism* to turn what its teams knew about their work into systems that changed how the work got done.

The intervention was organizational, not technological: a small **Forward-Deployed Engineering (FDE)** function — engineers embedded inside operational teams, chartered to ship working improvements against real transactions in weeks, within a governance spine that kept their autonomy auditable. Within twelve months, eleven of fourteen FDE builds reached production, inverting the historical pilot-to-production ratio. Just as important, the builds that shipped were revenue builds as much as cost builds: recovered rebate and billing leakage, faster quotes and approvals, fewer disputes delaying cash, and better-supported renewals.

02 The Current State: A Distance Problem

The organization ran four process towers — quoting and pricing, channel partner programs, global order flow, and billing and collections — alongside a three-region customer care function and enabling teams. Together they processed millions of high-volume, rules-heavy, document-intensive transactions a year. A diagnostic in early 2025 surfaced a consistent pattern:

- **Automation had plateaued.** A decade of RPA had harvested the easy wins. What remained involved judgment, exceptions, and unstructured documents — exactly where scripted bots break.
- **Pilots died in the handoff.** Ideas traveled from tower staff to business analysts to a central IT backlog. By the time anything was built, the requirement had drifted or the sponsor had moved on. Pilot-to-production conversion sat below 30%.
- **Knowledge and capability lived in different people.** Tower experts knew every exception path but could only express it as documents; central engineers could build but sat three layers from the transactions.
- **The real process was undocumented.** Process mining repeatedly showed actual flows — workarounds, side spreadsheets, approval shortcuts — bore little resemblance to the documented ones. Automation built to the official process failed at the undocumented steps.

The same pattern appears wherever operations run at scale. A manufacturer's warranty team keeps the real claim rules in a veteran's head; a telecom's provisioning fallout is fixed by side spreadsheets no system sees; a pharma company's chargeback analysts reconcile wholesaler data by hand because the documented process stopped matching reality years ago.

None of these was a technology gap. Each was a *distance* problem: the distance between the work and the people empowered to change it.

03 The Catalyst: Agentic AI Raised the Stakes

Through 2025–26, the technology acting on operations changed character — from bots that execute scripts to agents that reason toward a goal, such as validating an order against contract terms or resolving a support case end-to-end. This cut both ways. The opportunity grew, because the judgment-heavy work RPA never touched was precisely what agents could now address. But the failure mode worsened: an agent deployed onto a misunderstood process doesn't just break like a bot — it acts, confidently and wrongly.

Leadership concluded that **the constraint was no longer access to AI capability, but the ability to ground it in the granular reality of the organization's own processes — safely, and fast.** That is a deployment problem, and deployment problems are solved by people close to the work.

04 The Intervention: The Forward-Deployed Engineer

The model came from the software industry. A **Forward-Deployed Engineer** is a full-stack engineer stationed inside the environment where the work happens, shipping working software against real data and iterating daily with the people who do the work. Popularized by data-platform companies and spreading rapidly across AI firms, it rests on one observation: agentic systems succeed or fail on deployment context, not model quality.

The organization adapted it internally, embedding each FDE in one tower — with the order management team, the rebate analysts, the care floor — under three commitments:

1. **Ship against real transactions in weeks.** A first working improvement on live or shadowed data within 30 days, however narrow.
2. **Discover the process as it is, not as documented.** Every engagement began by mining evidence, shadowing operators, and mapping the minute steps and workarounds — because an agent built to the official process automates the fiction.
3. **Own the outcome, not the demo.** Success was measured on production adoption and the tower's own metrics, never pilots delivered.

> *An agent built to the official process automates the fiction. Discovery is not a preliminary step; it decides whether the agent works at all.*

EXHIBIT 1 **How the FDE differs from existing improvement roles**

Existing role	Why it hadn't solved the problem	What the FDE does differently
Central automation CoE	Builds from a backlog three layers from the work	Builds *in* the tower, on its priority, this sprint
External consultants	Knowledge exits with the engagement	Stays; knowledge compounds internally
Business analysts	Can't build; translation loses detail	Sees the work and writes the code
Citizen developers	Close to the work, but lack production rigor	Full engineering with production standards

Source: *Author analysis.*

05 Program Design

Scale and placement. Five FDEs for an organization of several thousand, one per major tower, reporting to the digital experience and insights function rather than central IT or any single tower — keeping them neutral, close to process-intelligence assets, and under a leader already accountable for transformation.

Operating rhythm. Each embed ran a repeating eight-week cycle: *discover* (weeks 1–2, map the real flow and pick one measurable slice), *build* (weeks 2–4, ship a first working version), *harden* (weeks 4–8, add guardrails and

integrate with systems of record), then *transfer* ownership to the tower — with updated role descriptions and training for the people whose work changed — and start the next slice.

Governance: autonomy inside a spine. A decision-gate model was agreed with finance, security, internal audit, and HR before the first build:

- **Deterministic guardrails** — margin floors, credit limits, payments, and ledger actions stayed hard-coded; AI could recommend but not execute past them.
- **Graduated autonomy** — shadow mode, then human-in-the-loop, then autonomy only for transaction classes with error rates below the human baseline.
- **Audit by default and a monthly gate review** — every agent decision logged; tower lead, FDE, and governance owner jointly decide what graduates; when a graduation changes how a role is performed, the tower's HR partner joins the decision.

Because the gates were pre-agreed, no build needed its own risk negotiation. Governance became an accelerant, not a brake.

06 What Changed: First-Year Outcomes

Eleven of fourteen builds reached production. Figures below are representative and directionally consistent with published industry benchmarks.

EXHIBIT 2 **First-year outcomes of production FDE builds**

Tower	FDE build	Outcome (year one)
Order management	Touchless order validation agent	Processing time down ~40% on covered order types; fallout cut by a third
Partner operations	Rebate and claim anomaly detection	Leakage on covered programs brought under 0.5% of program value
Quote-to-cash	Contract-to-invoice reconciliation assistant	Billing disputes down ~25%; reconciliation effort roughly halved
Customer care	Escalation summarization on case handoff	Handle time on escalated cases down ~35%
Quoting & pricing	Deal-desk triage agent	~70% of routine approvals pre-cleared within policy; queue time halved

Source: *Program data, year one. Figures are representative and directionally consistent with published industry benchmarks for each process area.*

Three second-order effects mattered as much as the metrics. The pilot-to-production ratio inverted, from under 30% to 11 of 14. Demand flipped: by mid-year, towers competed for FDE time and arrived with their own mined evidence. And the organization got smarter about what *not* to automate — three builds were deliberately stopped at the gate because discovery showed the process needed redesign, and leadership publicly counted these as successes.

07 Lessons Learned

Hire for translation, not just engineering. The best FDEs could sit with an analyst for a day, extract the unwritten rules, and keep the operator's trust while challenging the process. The failed first hire — technically excellent, floor-averse — taught this early. Section 08 sets out the full role profile.

Protect the role from both gravity wells. Central IT wanted to absorb FDEs into its backlog; towers wanted them as permanent support. Hard rules held the line: no ticket-queue work, and every build ships with a named tower owner.

Discovery is the undervalued deliverable. Every tower wanted to skip to the build; after one engagement, every tower cited the discovery map — how the work actually flowed — as worth paying for on its own. Timeboxing it to two weeks per slice prevented analysis paralysis.

Pre-agreed governance beats per-project approval. Negotiating the decision-gate model *once* was the single biggest speed unlock.

08 The Talent Side of the Model

The FDE model reads as an engineering program, but it rests on three talent decisions that operations leaders tend to discover late: who to hire, where to find them, and what happens to the roles around each build. Each is easier to get right with HR in the room from the start.

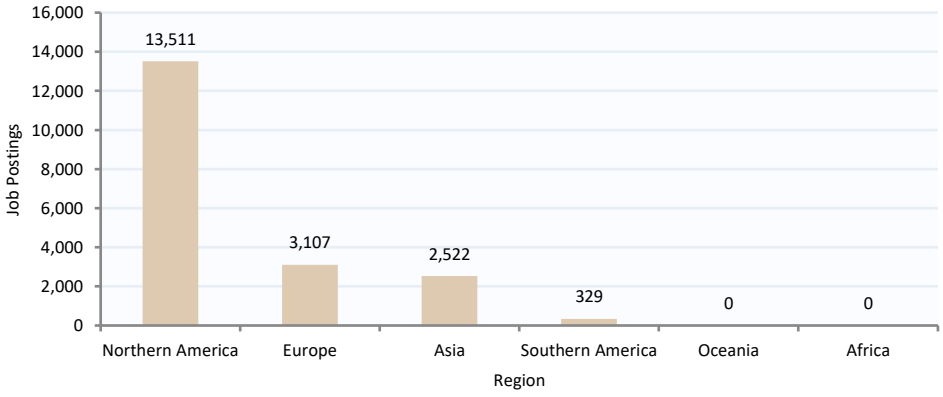
The role profile is rare by design. The FDE combines production engineering with the skills of a good analyst: shadowing operators, extracting unwritten rules, and challenging a process without losing the people who run it. Few candidates have both. Assessment should test the second half as seriously as the first; a floor-shadowing exercise reveals more than a coding screen.

Build and buy. External hires bring engineering depth but compete with AI companies for the same profile, at the same pay. Internal feeders — RPA developers, business analysts, and strong citizen developers — bring process knowledge and trust but need engineering uplift. A practical first cohort mixes the two.

The roles around the build change. When an order-validation agent cuts processing time by 40%, the order team's work shifts from keying and checking toward exception handling and customer-facing resolution. Deciding in advance whether released capacity is redeployed to revenue work, absorbed through attrition, or reinvested in new builds — and telling the teams — is what keeps operators willing to show the FDE their real process. Discovery depends on that trust.

Plan the career path before the first hire. Two years embedded in a tower builds rare cross-domain expertise, and without a visible next step these engineers leave. Clear leveling, a path into product, architecture, or transformation leadership, and pay benchmarked against the AI talent market rather than the internal IT scale keep the program's knowledge from walking out.

EXHIBIT 3 **The FDE talent market**

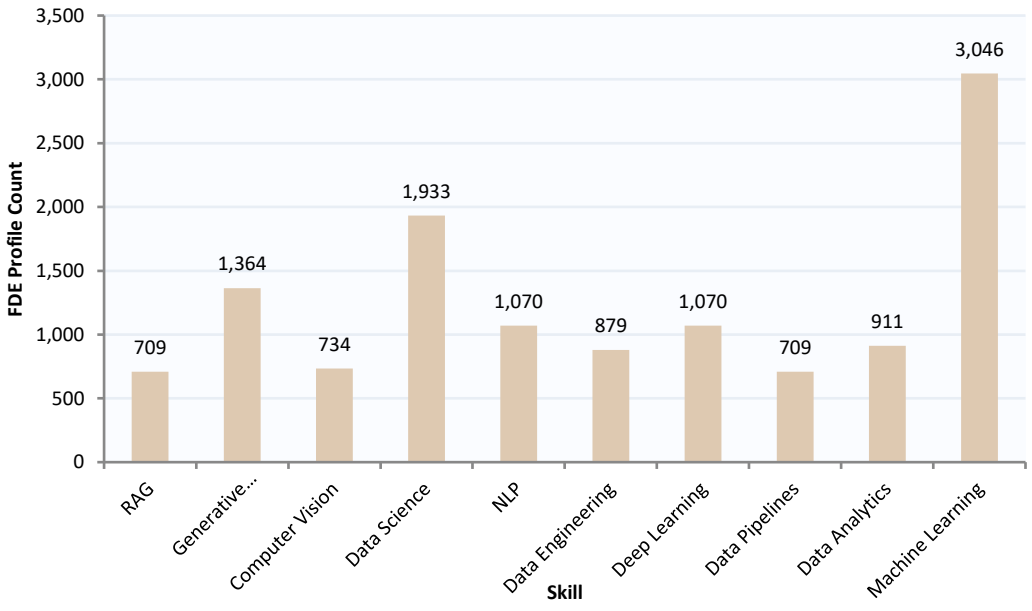
Indicator	Draup view
Demand growth; FDE Job Postings by Region, 2023–2026	

Indicator	Draup view

Top feeder roles: Top prior roles of current FDEs

Prior Role	Count
Software Engineer	561
Forward Deployed Engineer (lateral/prior FDE role)	190
Senior Software Engineer	154
Artificial Intelligence Engineer	105
Software Engineer Intern	94
Data Scientist	86
Data Engineer	66
Software Developer	63
Machine Learning Engineer	61
Solution Architect	59
Resident Solution Architect	50
Solution Engineer	42

Differentiating skills: FDE profile counts across top ten over-indexed skills



Source: Draup talent intelligence, September 2026

09 Applying the Model Across Industries

The mechanism is industry-neutral. What changes by sector is where the archetype towers live, which builds come first, and how high the autonomy ceiling sits — the level of autonomy regulation and risk will ultimately allow.

EXHIBIT 4 Where the model applies — archetype towers, first builds, and autonomy ceiling by industry

Industry	Where the archetype towers live	Representative first FDE builds	Autonomy ceiling
High tech & electronics	Complex-configuration quoting, order management, distributor programs, RMA and warranty, supply allocation	Configuration validation agent; ship-and-debit claim validation; RMA triage	High
Software & SaaS	Subscription and consumption billing, renewals, deal desk, channel, support	Renewal automation; usage-to-invoice reconciliation; deal-desk pre-clearance	High
Retail	Store and e-commerce order operations, vendor compliance, merchandising, returns	E-commerce order-exception triage; vendor chargeback validation; returns-fraud screening	High
Consumer goods (CPG)	Trade promotions, retailer deductions, replenishment, order-to-shelf	Deduction and claim validation; promotion-settlement reconciliation; order-exception triage	High
Industrial & manufacturing	Order-to-delivery, spare parts, warranty claims, dealer incentives, field service	Touchless PO intake; warranty-claim anomaly detection; spare-parts demand triage	High
Automotive	Dealer operations, warranty and recalls, supplier quality, parts logistics	Warranty-claim validation; supplier quality-deviation triage; recall-scope analysis	Medium–High
Telecom	Lead-to-activation, provisioning, billing disputes, retention and care	Provisioning-fallout prediction; billing-dispute triage; churn-save case assist	High
Logistics & transportation	Booking-to-billing, freight audit, shipment exceptions, customs	Freight-invoice audit; shipment-exception triage; customs document checks	High
Energy & utilities	Meter-to-cash, outage and field work, trade settlements	Billing-exception resolution; field work-order triage; settlement reconciliation	Medium
Financial services & insurance	KYC and onboarding, claims, loan operations, disputes	Claims document validation; KYC exception triage; dispute-response drafting	Low–Medium
Biopharma & life sciences	Chargebacks and rebates, regulatory documents, pharmacovigilance, trial operations	Chargeback validation; adverse-event intake triage; trial-site document checks	Low–Medium
Healthcare providers & payers	Revenue cycle, prior authorization, claims adjudication, patient access	Denial-root-cause triage; prior-auth packet preparation; claims-coding checks	Low–Medium
Public sector	Benefits and permit processing, grants, procurement, citizen services	Application completeness checks; case summarization for caseworkers	Low

Source: *Author analysis. Autonomy ceiling reflects the typical upper bound on agent autonomy given regulatory and risk constraints in each sector.*

Three sector differences matter; everything else transfers:

- **Regulation sets the autonomy ceiling, not the model.** In biopharma, healthcare, financial services, and public sector, many builds stay human-in-the-loop permanently — an adverse-event report, a prior authorization, or a benefits decision may never fully graduate. The decision-gate governance is identical; the graduation bar is higher. That is a setting, not a redesign.
- **Where process knowledge lives differs.** In high tech it sits in configuration rules and distributor data; in retail, in POS, order-management, and vendor EDI feeds; in consumer goods, in retailer deduction portals; on a factory floor, with operators and in MES data; in telecom, in OSS/BSS; in healthcare, across payer portals and EHR workflows. Discovery tooling changes accordingly; the discovery-first discipline does not.

- **Who the FDE must earn trust from differs.** Distributors in high tech, vendors and store operations in retail, retailers in consumer goods, dealer networks in automotive, clinicians in healthcare, and quality functions in life sciences all hold effective veto power. In these sectors the FDE's embedding extends to those stakeholders, and discovery covers their side of the handoff.

Where to start by sector. The best first builds share three traits everywhere: high transaction volume, a measurable leakage or cycle-time cost, and a clear human fallback. In practice that points to configuration and distributor-claim validation in high tech; order exceptions and vendor chargebacks in retail; deductions in consumer goods; warranty claims in manufacturing and automotive; fallout and dispute triage in telecom and utilities; freight audit in logistics; and document-completeness checks — where the agent prepares and a human decides — in regulated industries.

10 Conclusion: Replicating the Model

In the agentic era, the binding constraint on turning operations into a revenue lever is not AI capability but *deployment context* — honest, granular knowledge of how work actually flows, held by someone empowered to act on it. The FDE role manufactures that combination. For an organization in any industry, the replication checklist is short:

1. **Start with the diagnostic.** Confirm the failure mode is distance, not capability.
2. **Start small and embedded,** one FDE per major process area, reporting to a neutral transformation owner.
3. **Lead with timeboxed process discovery** and treat the map as a deliverable.
4. **Pre-negotiate the governance spine** once, before the first build.
5. **Bring HR in at design, not at handoff.** Agree the role profile, the sourcing mix, and what happens to changed roles before the first engineer is hired.
6. **Measure production adoption,** never demos.
7. **Celebrate the builds you stop.**

The economics are modest — a handful of senior engineers against an operating budget — and the risk is contained by design. What changes is the organization's metabolism: how fast an idea held by the people who do the work becomes a system that does it with them.

About This Brief

This brief presents an illustrative, composite case of a Forward-Deployed Engineering program in an enterprise operations organization. The organization, program details, and outcome figures are representative rather than audited results from a single named company; outcome figures are aligned with published industry benchmarks for each process area. The industry mapping reflects author analysis and should be validated against each organization's own process evidence.

About Draup

Draup is an AI-native work redesign platform. Its Etter platform maps how work is actually done — at the level of tasks, skills, and roles — and helps operations and HR leaders redesign it for agentic AI, supported by Draup's own forward-deployed teams.
https://www.draup.com, info@draup.com