

# On AI and Administration

---

Niket Desai

2026

*“At scale, the ledger becomes the work.”*

15 SECTIONS / ~10,000 WORDS



# 00

## Prologue: Ryan Leaves Law, and We Start Counting

*“At scale, the ledger becomes the work.”*

In 2016, a friend of mine, Ryan, told me he was considering leaving law. He had the credentials and the trajectory for a predictable career in litigation. The issue he kept coming back to was how his week actually worked.

He described to me two categories of effort. The first was the work people think they’re buying when they hire a serious firm — judgment, drafting, negotiation, client service, and decisions made under time pressure. The second was the work required to make the first category legible inside the institution — tracking activity, reconstructing it later, categorizing it into billable narratives, and producing artifacts that could survive client review.

His idea was to take advantage of advances in machine learning and apply them to legal timesheets. I told him not to leave law, for which he had spent seven years in higher education and a considerable sum of money to enter, and certainly not to start a company.

He did it anyway, of course.

In 2018, Ryan left, and we started what became Laurel (we were called Ping early on). We started with the premise that professional work was increasingly happening through digital tools — email, calendar, documents, messaging, and meetings. The accounting of that work, however, was manual, late, and unreliable. The institution was asking highly trained people to spend hours turning a week into paperwork — after the fact.

Professional services make this visible because time is priced directly. A timesheet is the firm’s representation of reality. It converts messy work into something the business can route, price, and defend. Once you see how an institution prices reality, you start noticing how many institutions survive on a similar ledger.

During an early pitch together at a Houston-based firm, a senior lawyer said he was buried in email. We had measured activity data that put a number on it. Many users were spending over three hours a day in email. We told them that if we could anchor digital work to a measurable surface — a ledger — we could change the workflow around it, including batching, routing, delegation, better norms, better tooling, and fewer forced handoffs.

Across pitches, deployments, and business reviews, we repeatedly noticed elite professional service firms with strong budgets and trained labor were paying a large tax in administrative reconstruction (Sunday night timesheet reconstruction, email archaeology, calendar backfills). We surmised the same pattern would exist in adjacent industries — and eventually in most institutions — but with less support and awareness.

I use “counting layer” as shorthand for that tax. It is the work required to make work legible to the system so the system can subsequently route, price, and justify it. The counting layer grows with scale. Eventually, it consumes the capacity of the system it supports.

This essay started as observations from building workflow automation software. It became an argument about how institutions work, why they slow down, and what changes when machines can carry part of the load. In this essay, I use “sophisticated compute” to mean machine learning deployed into workflows at scale. It argues a stacking thesis.

*Work is made of transactions. Transactions run through workflows. Workflows survive because they turn messy reality into something routable, repeatable, and defensible — but that translation requires a counting layer. The forms, backfills, and reconstruction that make work legible to the institution. As scale rises, the counting layer grows until it steals capacity from the work itself. The promise of sophisticated compute isn’t magic; it’s the boring ability to carry part of that counting layer without hiring humans in proportion. If it works, service expands. If it fails, it fails at the workflow level.*

The rest of the essay is an attempt to name that counting layer precisely, show where it hides, and describe what “carrying it” looks like without breaking the workflows institutions rely upon.

CONTINUES → 01

It starts with a lesson we learned early. Section 01 explains why most AI conversations miss this. Capability is easy to demo. Integration is what institutions buy.

# 01 Talk Is Cheap

*“Capability is cheap. Integration is expensive.”*

In the early 2020s, the public met AI directly through language models. You type, it responds, and the response is interesting enough to feel like a small party trick. That entry point turned a field most people couldn't touch into something anyone could.

That same entry point distorted the conversation. It trained people to judge AI by what it could demo. Generative LLMs — and diffusion models behind much of what's produced online — pull attention toward instant feedback. Institutions, as reflected in renewal decisions, are not moved by parlor tricks for long. They purchase structural outcomes they can depend upon.

We learned this lesson early at Laurel. Our first instinct was to lead with capability. Automation, intelligence, and the idea that a machine could “understand” work. Buyers were curious, sometimes excited, and then immediately practical. They wanted better revenue realization, an airtight compliance posture, and fewer workflow failure modes, especially the kind that ends with IT holding the bag. The buyer team (procurement, risk, and IT) kept returning to the same question. What happens when this is wrong, and who owns the downside? Ryan carried more than his share of those conversations while we were earning the right to claim reliability and repeatability.

The same thing happens everywhere. In any mature workflow, AI is rarely the product. It's buried inside billing, claims, and scheduling — inside the systems people actually buy. If the outcome isn't routable, repeatable, and defensible, capability doesn't matter.

LLMs make the trade visible. Their advantage is flexible input. You can throw messy content at them and still get something usable. Their disadvantage is output volatility. In a high-volume, low-risk system, generation volatility becomes the work. Prompting, tuning, evaluation, routing, fallbacks, refusals, monitoring — the operational effort is spent turning volatility into something dependable.

If the workflow is an artery — think revenue, trust, compliance — shallow rewiring isn't an upgrade. It's a risk, even when buyers don't spell that out cleanly in an RFP.

Embedded workflows don't change because a better tool shows up. They change when standing still becomes more dangerous than moving, or when someone inside the system can map the workflow comprehensively

enough that not rewriting becomes indefensible.

A managing partner at a London firm was one of the first people we met who behaved like an operator in this sense. He wasn't asking what our AI could do in a vacuum when he chose to back Laurel early and become one of its first customers. He was asking what would change if work could be measured and acted on, in a market he expected to move toward client-pressured fixed fees and away from unconstrained hours. While being a design partner, he and his team kept us focused on outcomes. If anything, they made the real adoption constraint explicit. Alignment inside the institution matters as much as the technology.

Buyers are skeptical of AI add-ons sold as systems; plenty of teams have been burned by thin integrations that demo well but don't survive deployment. Incumbents compound the problem. They market capabilities they rarely deliver — sometimes because the tech won't support them, more often because the organization won't tolerate the required workflow self-rewrite. Their safer differentiation is bundling. Sell suites to reduce procurement friction and deepen lock-in while leaving core products largely unchanged. Meanwhile, the work that actually drives adoption — deployment, integration, security review, change management, ROI validation — often takes longer than a startup runway allows.

CONTINUES → 02

AI adoption fails or succeeds at the workflow level. Section 02 defines what a workflow is, why it becomes embedded, and why many “AI integration” attempts stay shallow by default.

# 02 Workflows and Routing

*“Every workflow was a choice before it became the constraint.”*

The world has always been a routing problem. People, goods, messages, money, authority. First, you choose how to move something, and then you standardize the container. Then, you scale the trade.

A workflow is routing made enduring. It takes messy reality and turns it into something the system can accept, whether a payment, a record, a decision, a delivered service, or a completed case. It's the difference between “we did the work” and “the work exists in a form that the institution recognizes.”

Modern life increases routing demand because the population has grown by billions since the late 20th century, and connectivity has become ubiquitous. Along

with this came more digital surfaces per person, including email, chat, calendars, documents, tickets, CRMs, payments, identity systems, and an endless assortment of internal tools to handle problems mildly differently. More surfaces mean more transactions.

Workflows around a given transaction become embedded standards because they accumulate a powerful ecosystem around them, including training, job roles, billing conventions, compliance standards, reporting, audit posture, procurement logic, informal habits, and even identity. These standards endure because they coordinate volume.

In professional services, the time ledger hardened in stages. In the 1910s, Reginald Heber Smith pushed modern law-firm management discipline — accounting, budgets, profit distribution, timesheets, and what became the billable hour — then institutionalized those methods during his long run as managing partner at Hale and Dorr (1919–1956). In the early 20th century, Paul Cravath built what became known as the Cravath System — the recruiting, training, and promotion machinery that turned elite firms into scalable institutions. Then, in the early 1970s, hourly billing became the de facto pricing standard for large-firm practice — not because it was beautiful, but because it was routable, defensible, and easy to audit at volume. Once an industry builds around artifacts like these, you don't "replace" them when new technology comes along. Many digital advancements from the 1990s onward dematerialized the paper, supercharged the tracking, and entrenched the ledger even more deeply.

A shipping container isn't a box. It's ports, cranes, customs systems, insurance norms, standardized dimensions, schedules, and contracts. A better container design doesn't matter unless the rest of the world agrees to handle it. Workflow change works similarly. You're not changing the step; you're changing what depends on the step.

This is why much of AI integration is shallow by default. The easiest play is to attach a new tool to an existing workflow and call it modernization. For some businesses, the appearance of modernization matters more than modernization itself. That can be useful at the margins. As an entrepreneur, these maneuvers are frustrating because the real opportunity gets missed. Rewriting the workflow itself so the institution routes work differently, counts it differently, and pays for it differently.

It's also why disruption can be a misnomer in the way it's often used. In the last wave, technologists who understood software but not industry workflows digitized those workflows from the outside — and many built

enormously valuable companies doing it. That wasn't disruption so much as fighting with better weapons. The industry experts had the knowledge but not the tools; the technologists had the tools but not the knowledge. Now, as sophisticated compute becomes broadly available, that asymmetry is closing. A lot of advantages can move toward industry insiders — the people who understand the incentives, the failure modes, the compliance constraints, and the murkier places where systems get gamed.

Think of the propeller. New manufacturing methods make new forms possible, and new forms change what performance looks like. In this essay, sophisticated compute plays the "new manufacturing" role by making work measurable at depth and scale — raw material for workflow redesign. I'll return to the mechanics in Section 05.

A new propeller shape doesn't get bolted onto a plane and waved through. It has to cooperate with certification, protocols, maintenance, training, liability, and procurement. Digital systems have a similar reality, even if it doesn't look that way from the outside. Tools change fast. Workflows don't. They're held in place by the larger ecosystem, which depends on them.

CONTINUES → 03

Workflows are routing standards that support and drive transaction volume. As volume increases, the support cost grows — the counting layer — and, over time, it eats at the capacity of the system. Section 03 names what that cost becomes when it hardens into a permanent layer.

## 03 On Administration

*"Administration eventually scales faster than value does."*

When I was younger — coming out of university — it was easy to treat administration as a moral failure. The campus version of the world hides a lot of the necessary machinery to focus on ideas and subject matter. We're implicitly led to believe administration is incompetence. You believe in "flat organizations" (or what the youth call group projects) and "no process" at face value before attempting to build anything that has to survive volume and time.

At Laurel, we saw this from the inside. The firms we worked with weren't poorly run. They were buried in the machinery required to keep work legible at volume. The pattern wasn't unique to law; it was administration doing what administration does.

Administration exists because trust doesn't scale. In a small team, you can verify directly — you see the work,

you know the people, you trust the outcome. Charlie Munger put it well, that the highest form a civilization can reach is “a seamless, non-bureaucratic web of deserved trust.” But trust requires consistency, and consistency at volume is itself a counting problem. When an entity outgrows the circle where trust is direct, it builds a ledger — routing, verification, records, accountability. The system that maintains that ledger is administration. It lets workflows run repeatedly without renegotiating reality for each transaction.

Bureaucracy is a subtype of that layer. It’s what administration looks like when it becomes self-protective, slow to update, and more committed to procedure than outcomes. It’s commonly veiled as security or self-preservation (colloquially: CYA).

Administration also creates abstraction. As the counting layer grows, work gets pushed further from end users and further from the heart of value creation. The institution becomes better at producing institutional artifacts and worse at producing service.

None of this happens overnight, and that’s why the creep can look like progress. A proxy starts as a measurement. Then it becomes a target. Then it becomes identity. The people closest to end users feel the distance first. Founders feel the creep early because their original vision remains in their head as a counterweight to the counters.

I respect why administration exists. I’m frustrated by what it does at scale. A lot of systems have held together longer than they had any right to. A common example in government contexts is “emergency powers” that outlive the emergency. Cincinnatus would undoubtedly be disappointed.<sup>[CIN]</sup>

The pattern has a human consequence. Some people use administration as a tool. Others have made it their whole job — mistaking gatekeeping for competence, converting process into power. In large institutions, that dynamic multiplies the counting layer without multiplying service.

You can see the pressure in who leaves. In many fields, the highest-skilled operators migrate toward roles with less administrative reconstruction — inside roles, narrower scopes, advisory positions, or, in exceptional cases, starting something smaller where the work stays closer to the value (big law to in-house; banking to buy-side; public service to advisory). We build entities that generate administrative overhead and then drag our experts into it.

Efficacy drops. Supply drops. Demand keeps rising.

While administration keeps systems coherent, it starts consuming the capacity it was meant to build and serve. As transaction volume rises faster than teams can manually reconstruct, verify, and record it, sophisticated compute becomes a viable way to carry the counting layer without hiring an ever-growing number of humans to do it. Section 04 puts numbers on why.

[CIN] Lucius Quinctius Cincinnatus is a figure of early Republican Rome. Tradition holds that in 458 BCE he was appointed dictator during a crisis, won quickly, then resigned and returned to his farm — cited as a model of civic virtue and restraint in power.

## 04 Coordination Gravity in Entities

“Coordination is the hidden tax.”

Most people think about scale in terms of outcomes like more customers, more revenue, more coverage, more “impact.” The harder part is the input side. When an entity grows, the number of people who must coordinate grows, and the coordination burden grows faster than headcount suggests.

A simple model is nodes and edges. Humans are nodes, and every relationship that requires alignment is an edge. With one person, there are no edges — just a single operator connecting all the dots. Alignment is simpler because it lives in one mind, and scalability is limited by one set of hours. With two people, there’s one edge. With four, there are six. In a fully connected group, the number of edges is:

$$A(h) = h(h - 1) / 2$$

COORDINATION EDGES

This shows why coordination pressure grows on the order of  $h^2$ . When I first encountered it in university, I didn’t find it prophetic of what organizations would attempt to solve with hierarchy, roles, and process — and, to their credit, those solutions reduce the \*effective\* edge count even if the underlying math still looms.

A shocking thing happens when you look at coordination in magnitudes instead of headcount (these are approximate).

A(h) EDGES	h ≈ PEOPLE	ORDER
10	5	 10 <sup>1</sup>
100	15	 10 <sup>2</sup>
1,000	45	 10 <sup>3</sup>
10,000	142	 10 <sup>4</sup>
100,000	448	 10 <sup>5</sup>

What we casually call a “small company” can experience a 10,000-edge coordination problem. Increasing from 50 → 500 people does not introduce a 10× complexity, but a 100× one. That’s before you add multiple channels, multiple time zones, multiple products, multiple customer segments, and the growing requirement that decisions be documented for handoffs, audits, and reuse.

When I was Chief of Staff at Flipkart, we hit a growth regime where the company was adding on the order of tens of millions in GMV and hundreds of hires per week, sustained for long stretches. While much of that was absorbed in fulfillment and delivery, the coordination complexity the team was managing was extraordinary. It also put into context the massive entry classes I later saw at Google. Some of the old guard were starting to exit as a once-small, once-legible company became unrecognizable to them after years of compounding growth.

“Fewer layers” becomes the cyclical reaction when management teams “right-size” organizations that have gone awry. Beyond founders, few operators have the courage to cut deeper than they expand. Markets reward growth; they rarely reward restraint.

When teams are small, more people encounter the market and the customer directly. Inputs aren’t filtered into generalized reports with long feedback loops. You make better decisions because you see the consequences earlier. When teams are large, a growing portion of the work becomes the work of staying aligned (or not creating warring parties).

Entities try to manage this in predictable ways through introducing structure. They insert management layers. They formalize the process. They build compliance functions, planning functions, cross functional coordination roles, and review rituals. They don’t do this because they’re evil; it is a response to expanding value through volume.

Proxies replace reality because proxies are easier to share than reality. A weekly report is easier to route than a thousand customer conversations. A quarterly plan is easier to defend than admitting, in real time, that you changed your mind three times. And long-range board plans often harden into commitments that outlive the facts that produced them. The artifacts become legible — and soon they become the target. When the artifact becomes identity, you start optimizing for the artifact.

Large entities slow down even when they’re full of smart people. The entity doesn’t lack intelligence. It lacks the bandwidth to coordinate that intelligence without paying a high tax in administration and counting.

I think of it like a rocket carrying its own fuel. As the rocket gets larger, more of what it carries is fuel for the fuel. The payload becomes the minority of the mass. The rocket still flies; it just becomes expensive to steer. Markets sometimes reward this tradeoff because scale can buy a durable advantage. I’m not sure many organizations can realistically aim for that outcome.

Coordination gravity is my shorthand for this. The force that pulls entities toward administration as transaction volume and internal complexity rise. It’s the same gravity that makes organizations reach for process and artifacts.

In earlier eras, scaling meant people and paperwork. Digitization turned artifacts into software. Globalization turned software into a 24/7 relay across time zones. Centralization followed because the entities best at running that relay got better at collating data, standardizing workflows, and concentrating expertise — advantages that compound and are hard to catch. That system worked until coordination and counting became the bottleneck.

CONTINUES → 05

Sophisticated compute is a new scaling layer that can carry part of that load — and Section 05 explains what it can (and can’t) replace.

## 05 Enter Sophisticated Compute

*“Anyone can rent the machine — few know what to build with it.”*

Sophisticated compute has been invented and advanced for decades by extraordinary scientists and engineers. What changed in the 2010s and 2020s is general availability. Frontier capability became purchasable infrastructure.

For much of the late-20th and early-21st century, the high end of machine learning was concentrated inside entities with outsized budgets — top research teams, proprietary infrastructure, serious security organizations, and the ability to run long experiments without the company dying in the process. A professional services IT team wasn’t going to compete with Google’s infrastructure or security posture. They were busy keeping the lights on.

Now a number of those capabilities are purchasable as infrastructure. They show up as APIs and managed services. This matters because it changes who gets to participate and deploy these technologies. More importantly, it changes who gets to redesign workflows.

I use “sophisticated compute” throughout this essay. “AI” has become shorthand for everything. “Compute”

keeps the focus grounded in deployment, capacity, and constraints. It also avoids treating language models as the entire field. LLMs matter because they're visible and they're a major interface — but they are not the totality.

Availability changes the existing advantage structure. When capabilities are rare, the advantage sits with the entities that can build them. When capabilities become common, the advantage moves toward those who can apply them at points of leverage. In this wave, workflow understanding becomes leverage.

I've written elsewhere about the operator and the agent, but the short version here is simple. Sophisticated compute is most powerful when it amplifies the people who understand the work. In this wave, deployers inside industries matter as much as model builders. If we don't create a generation of deployers, we don't get modernization; we get additional vendor dependence.

In particular, operational data in these industries can make the difference. Public data is table stakes and is now incorporated into baseline services. When you ask Google a query, you're no longer asking "the web" and are instead receiving Google's interpretation of the web. That interpretation is the product.

The same pattern will emerge as institutional workflows become better mapped and interpreted. The most valuable data won't be on the public web; it's operational exhaust generated by workflows — who did what, when, with what dependencies, what changed hands, what failed, what got reworked, what was approved, and what was billed. Much of it is trapped inside tools that were never designed to be analyzed as a single surface — so the competitive advantage is not "having a model," but earning a permissioned, integrated view of the workflow.

Sophisticated compute elevates beyond "AI added to a workflow." It is a new manufacturing method that changes what designs are feasible — and it can scale existing designs beyond what manual operations can reliably support.

CONTINUES → 06

But moving beyond "theoretically possible" into production requires deploying in a way that survives the triad of revenue, trust, and compliance. It has to get paid for, it has to be dependable enough to earn repeat usage, and it has to satisfy the institution's rules. Section 06 breaks down how that triad shapes what actually gets deployed.

## 06 How Automation Enters Work

*"Automation is a routing decision."*

Automation is not new, and it is not a monolith. It enters organizations as lanes. A track that can be automated at high volume. A hybrid track where automation assists, but humans stay in the loop. A human track for exceptions, judgment calls, and outcomes where the cost of being wrong is unacceptable.

Consider invoice intake for accounts payable. Clean invoices with matching purchase orders can run fully automated. Invoices with missing fields, mismatched line items, or ambiguous vendor names move to the hybrid lane — compute proposes a match, and a human confirms or escalates. High-dollar invoices, duplicate signals, and anything with fraud exposure stays human-led.

The proportions shift over time as organizations tune the mix of people, process, and technology.

To build reliable lanes, you have to deeply map the workflow. You have to know what the transaction is, what the institution treats as a valid outcome, what the failure modes are, and who owns the downside when it goes wrong. With that understanding, it becomes clearer which transactions can be grouped, deduplicated, routed automatically, require escalation, and need audit trails.

If that work is skipped, you get what the market is overwhelmed with right now. Incredible demos attached to shallow integration points. It looks like modernization until the first serious exception shows up. Then, the institution does what institutions do. It routes the exception to a human, and it routes the tool to the graveyard.

Tools and their integration paths implicitly choose a point on what I'll call the deployment triad. Latency is the cost and speed of producing an output at volume. Accuracy is whether the output matches what the workflow actually needs, in context. Safety is whether the system can be deployed without unacceptable harm, including harm that shows up as liability, reputational damage, or regulatory action.

Few systems clear a high bar on all three dimensions, and setting the minimum acceptable bar is where the pain lives. You either accept latency costs, accept residual error, accept a narrower deployment surface, or accept a heavier reliability layer. And the client always finds out which choice you made — and shipped.

LLMs make the trade visible. Flexible input expands what you can route. Volatile output expands what you must contain. In practice, that pushes effort into the reliability layer — evaluation, routing, fallbacks, refusals, monitoring — so the institution can defend the outcome.

The easiest deployments are the ones you can unwind, but they are also the ones most exposed to competition. The hardest deployments sit inside low-reversibility workflows — places where once you ship a change, everything downstream assumes it is true. Reversibility is the property that determines risk. How hard it is to roll a deployment back once real money, trust, and compliance are resting on it.

You can see two different deployment philosophies. Certain systems are built like aviation: constrain the environment, harden the workflow, expand slowly. Others are built like consumer software: ship early, let the market select the use cases, and build containment after the fact. Neither approach is universally correct. They map to different downside profiles and risk tolerances.

With AI, people fear lane models because the hybrid lane can feel like training their replacement. That fear is rational inside systems where the institution treats labor as a cost center, and the worker treats the job as an identity. In practice, the threat isn't just income; it's the income/identity/comfort bundle people build their lives around. It's difficult for organizations to thank people for their hard work and tell them it was mostly counting. The only durable answer is to make lanes explicit and design for upward movement. Fewer humans doing rote counting, more humans doing judgment, collaboration, and the work that benefits from human loops.

CONTINUES → 07

If this lane structure is obvious, why isn't everyone doing it? Because making lanes real — automated where safe, hybrid where useful, human where necessary — runs into deployment constraints. Procurement, compliance, security posture, and the incentive structure of incumbent players. Section 07 explains why those constraints — not capability — decide what gets adopted.

## 07 The Gates and the Rake

*“Gates are rational, rakes are profitable.”*

Lane models run into a separate constraint that has nothing to do with model capability. Institutions have gates, and gates have (for the foreseeable future, human) owners.

Many gates are legitimate. If a workflow is a backbone artery — revenue, trust, compliance — the buyer team inherits the downside. Whether they understand the institutional-level implication or not, they understand that mature workflows come with mature friction learned through pain, including security review, privacy posture, audit trail, continuity planning, escalation paths, and a penchant, possibly pleasure, for worst-case thinking.

We ran into this early, and to this day, at Laurel. RFPs, security questionnaires, SOC 2, DPAs, pen tests, vendor assessments, and other equivalents when trying to modify one of, if not the most critical, workflows at firms that drive money and policy. Sometimes there were multiple versions of the same gate, each run by different teams and with different incentives. The process was long and convoluted, but not irrational. We heard enough stories of tools and implementations going wrong to understand why self-protection was rational.

Within many firms, the IT and security teams maintain systems they didn't design or build and don't have the mandate to replace. Their world is asymmetric. Upside is rarely rewarded, downside is punished immediately (and in the case of lawyers, rather scathingly). In that setup, opportunity looks speculative; failure looks career-ending. It is hard to sell “innovation” to someone whose job is to prevent surprises.

Over time, gates turn into rituals. When a decision is made and then handed over to the implementing team, a checklist transforms into a ledger of plausible deniability. Someone “owns” the paperwork; it gets filled out, and the institution can say it did its due diligence. The gate stops functioning as a filter for quality and starts functioning as a shield (sometimes a sword) for the status quo. And that's where the rake shows up.

Incumbents in mature industries have learned to brand legacy postures as “safety”—on-prem requirements, bespoke integrations, proprietary formats, suite purchasing, single-vendor procurement logic.

The practical effect is to keep workflow standards embedded and to keep rewrite pressure low. Innovation is cut down in favor of product suites that simplify procurement while leaving core workflows untouched. For startups, all the value is in workflow modification. For incumbents, all the value is in workflow maintenance.

COVID exposed something adjacent. Between 2015–2019, Ryan and I were told that cloud-based systems handling core financial workflows at law firms would not be taken seriously. When remote work became a continuity requirement, objections softened and shifted

into deeper reviews (the sales cycle remained painstakingly glacial). On-prem had been treated as the default posture in many industries until continuity forced a comparison. They didn't change out of faith in better tools; they changed because the risk moved, and standing still became indefensible.

The current AI market has further tightened the gates. Thin tools can be assembled quickly and sold like complete systems. Plenty of buyer teams have been burned. Their response has been to protract the sales cycle with longer due diligence periods, more requirements, and more paperwork — even when labeled as design partnerships. The effort in workflow change — deployment, integration, security review, change management, ROI validation — keeps expanding, while the average startup runway does not. There will be an entire generation of startups that fail because of disasters from the first wave of AI workflow builders.

The pressure from financial markets and opportunists to drive AI adoption can lead to discourse that feels detached from what is happening inside institutions. In public, capability is the headline. In private, the story is whether the workflow can survive and who carries liability when something inevitably goes wrong.

CONTINUES → 08

Gates do not stop modernization, but they shape where it happens. As coordination gravity increases, large institutions avoid rewriting durable workflows because the cost and downside are too high. So, modernization often happens adjacent to them. New, smaller entities built to meet the same security, compliance, and continuity requirements without inheriting the accumulated coordination and administrative weight of the incumbent stack. Section 08 describes what those adjacent institutions look like — and why they can move faster without breaking the rules.

## 08 The Many-Entity Market

*“Small teams will do outsized work.”*

Coordination gravity is a superlinear tax. For institutions, the question becomes whether it remains a tax paid in human headcount, or whether sophisticated compute can carry part of the counting and routing load without adding humans in proportion. If it can, the shape of the firm must change.

The first-order effect is that the institution becomes easier to steer. Teams that used to require dozens or hundreds of people to operate a workflow can shrink, while output stays constant or increases. Administration per transaction drops. The distance between builder and customer shrinks. The many-entity structure doesn't only describe markets; it describes what happens inside firms.

A CFO builds a complete RevOps function with three people and heavy automation instead of a fifteen-person department. The logic is the same at every scale.

The second-order effect is market structure. If small teams can deliver credible outcomes in narrow slices of work, you don't get a single winner swallowing adjacent industries by suite expansion. You get many entities competing inside sub-verticals — by industry, geography, language, regulatory regime, risk tier, customer segment, or even a single workflow step. In the mid-2020s, the U.S. started treating data, compute, and cross-border workflow dependencies as instruments of statecraft; Europe's response has been sovereignty — rebuilding critical workflows (and the corporate wrappers around them) inside European jurisdictions and compliance regimes.

You can see what I mean in KYC/AML. One entity specializes in identity and document verification for a single jurisdiction. Another owns sanctions screening and adverse media, tuned to a risk tier. Another is case management and audit trails for the exceptions. None of these needs to own “the whole suite” to be essential — and once the work is narrow enough, small teams can credibly compete.

A world with more specialized entities will deliver better services to more people because the cost basis drops and the service surface expands. When the workflow is safe to run at volume, marginal cost declines, and growth is made up in additional volume, ideally making the workflow safer to run. The lower cost basis and wider availability mean the average person can access services that were previously reserved for a narrow set of buyers.

This shows up differently across economies. In the U.S., labor is expensive, so many services are rationed by labor cost. In places like India, labor is cheaper, but caseloads are massive, and that scale bends or breaks the human model in a different way. Sophisticated compute lowers human minutes per transaction while making new transaction volume possible.

The best outcome of “AI everywhere” is more transactions that produce value — think reviews, more follow-ups, more compliance checks, more casework, more interpretation, more education — without adding a matching or greater human administrative burden.

Many entities create a second coordination problem, not inside a firm, but between firms. At fifty entities, you can still coordinate with a handful of shared standards. At five hundred, the edge count explodes, and the “integration layer” becomes its own counting layer. Identity, permissioning, audit logs, and data portability

become the new ports and cranes — the boring interfaces that make specialization composable instead of brittle. Without shared interfaces, specialization stops being a market and turns into a pile of bespoke integrations that recreate the counting layer between firms.

The many-entity economy is not guaranteed. It is a design direction that requires builders who have intimate knowledge and respect for industry workflows, who can rewrite them from the inside. For each seasoned software engineer or entrepreneur, there are multiples of experts who are limited by their technical proficiency; that limitation will drop as interfaces improve and deployment becomes more standardized.

CONTINUES → 09

As these industry experts rebuild their workflows from within, as opposed to awaiting a technologist outsider, there will be significant displacement. Section 09 is where that bill gets paid.

## 09 Transition Costs and the Labor Reallocation

*“Nobody thanks you for a decade of counting.”*

Displacement is the visible cost of this transition, and it is not evenly distributed.

If you can describe a role as repeatable routing and counting — work that does not require meaningful creativity or collaboration — then sophisticated compute will absorb pieces of it. This happens because those roles exist to compensate for limits in tooling and coordination.

The old sequence from the research and development world was people → process → technology. In certain domains, the sequence flips technology → process → people. Start with what can be automated safely under load, design the process around it, and concentrate humans where the workflow requires judgment, discretion, and human loops.

This transition hits income, identity, and comfort in the same motion. People are rightfully scared about money and being able to support themselves and their families. People are scared of narrative. People are scared because their work is proof of worth in certain societal constructs. They might even defend the work they do not enjoy when feeling these emotions. In professional services, we see people terrified that parts of the value chain are less differentiated than they believed, while simultaneously resenting the administrative overhead they never wanted to do.

The timing is cruel for people late in their careers. If you are late in your career and an industry rewrite lands

inside your working lifetime, it feels unfair. The institution told you a story about long-term stability — and then renegotiated it. No essay I write fixes that; the practical question is what a society does in response.

My view is that transaction volume tends to rise geometrically (people x digital surfaces x connectivity = more transactions). The future requires automation and compute to support it, even if we dislike the transition. Slowing advancement does not remove or suppress demand; it pushes demand further into overload, where burnout, shortages, and systems become brittle under volume.

We can return to our lane models as the mechanism. The automated lane absorbs repeatable counting. The hybrid lane absorbs assisted work while the system learns. The human lane concentrates on exceptions, contested judgment, and the parts of work that cannot tolerate error. As the mapping improves and the thresholds get tuned, more transaction types move safely into higher-volume lanes. Humans spend fewer minutes on administrative steps per transaction and more time on judgment calls and continuous improvement.

That transition works at the system level. Retraining individual workers is harder — in digital and counting-heavy industries, the purpose of these roles is the counting layer that just got compressed. There is no adjacent rung to step onto. That’s why compensation inside work — and support outside it — are a vital bridge. New categories of work have always emerged after major transitions, but they don’t arrive on the same timeline as the displacement.

Different countries will implement varying levels of worker protection and slower adoption, but many companies have shown a brutal pattern of organizational cuts made with little protection offered to outgoing employees.

To the entrepreneurs out there — creating jobs has never been more important.

CONTINUES → 10

The worst thing we can do is pretend everything will be fine, because we will land in the worst of both worlds. This transition cannot be managed purely inside the market. Governance will end up carrying the pressure through a few blunt levers such as income bridges, retraining tied to workflows, and deployment regimes that set safety rails. Government runs the largest workflows, and when those workflows fail, the cost accumulates inside the system as backlog, litigation, and backlash. Section 10 is where that becomes unavoidable.

# 10

## Government as a Workflow Operator

*“The state runs on workflows; citizens live inside them.”*

Government is not a metaphorical institution. It is a scaled operator of workflows. It routes money, authority, benefits, permits, enforcement, and claims. And it operates under constraints that private entities don't. Citizens can't meaningfully switch providers, and the state can compel participation. When government workflows fail, the cost doesn't resolve through churn. It accumulates as backlog, litigation, and backlash.

What citizens experience as “the government” is a set of transactions moving through hardened workflows. Taxes, benefits, licenses, immigration adjudication, permitting, procurement, court scheduling, case management, emergency response, and more. While administrations come and go, the backlog (and backlash) of government transactions across their varied services remains.

I am a proponent of AI use in government contexts as a key mechanism to improve efficiency and, more importantly, the efficacy of the services and programs citizens have purchased. Government modernization perpetually lives in procurement and audits. And while budgets increase at every level, from municipal to federal government, it is not clear that we are receiving a better marginal return. Much of the debate on the bounds of government necessity is predicated on an incomplete understanding of the service workflow.

In Section 07, I described gates as rational. Security review, audit trail, continuity planning, escalation paths. Government has all those gates, plus coercive power and public scrutiny. Procurement becomes governance by other means. The rules that decide what can be bought become the rules that decide what can change, and overloaded diligence can quietly guarantee ossification.

When people talk about modernizing government, the temptation is to reach for a single move. “Replace the old system with a better one.” Look at how many reformers are elected on “replace the old system,” and discover that durable change dies inside the existing workflow stack. The government system is a dense set of workflows that have been patched for decades, with multiple ledgers and owners of risk.

If you want a government to run better, the first act is reliable, transparent measurement. Choose the transactions and map the routing to reveal the counting layer. Define the human minutes per transaction. Then

decide what gets an automated lane, what stays hybrid, and what must remain human because the public cost of a wrong answer — appeals, rights challenges, court involvement, and backlash — is higher than the throughput cost of a slow one.

This is why I'm skeptical of opaque “efficiency” pushes that don't start with counting. When rhetoric gets ahead of measurement, you get an argument about intent instead of a record of outcomes. And personally, I struggle with a society that remembers the last thing spoken more than the record of what was done.

My argument can be exploited by the new age of vendors.

The U.S. government has increasingly moved toward large enterprise contract vehicles for cloud, software, and data infrastructure with a small set of vendors. At the platform layer, the Pentagon's JWCC cloud contract is explicitly a multi-vendor attempt to standardize procurement without standardizing on a single provider. At the application-and-data layer, the U.S. Army's Palantir Enterprise Agreement consolidates prior contracts into a long-term framework with a large ceiling.<sup>[ARMY]</sup>

Those deals can be a much-needed upgrade. They may also become the next generation of workflow standards. The vendor ends up defining the interface and the data schema, and therefore the ledger, routing, and permissible kinds of onward change.

Competition within government matters the same way it matters in any other industry. It fends off the next ossification. Our government(s) do important work at scale that other markets will not, and for that they deserve access to first-class tooling and capabilities.

There is a reason this is hard. A permit workflow is annoying. A benefits workflow could be radioactive. A justice workflow could be existential. In every one of those cases, the counting layer steals capacity from the frontline. How do we reconcile a police officer's choice between being available to respond to a call for service and the administrative reconstruction — reports, evidence chain of custody, and documentation that exists to ensure the system can defend itself later?

Sophisticated compute can help carry the counting layers in government so that humans spend more time where the work is human — judgment, discretion, fairness, and compassion. Generations of governed citizens live and die within these workflows.

CONTINUES → 11

Other nations notice. Section 11 frames the emerging competition between deployment regimes.

[ARMY] The U.S. Army describes its Palantir Enterprise Agreement as a volume-discount enterprise framework with a not-to-exceed cap of up to \$10B over up to 10 years (a ceiling, not obligated spend).

# 11 Competing Geopolitical Deployment Regimes

*“A society exports its values through its infrastructure.”*

The way a government carries its counting layer is not a technical choice. It is a statement about what kind of society it intends to be. The workflows that route money, benefits, permits, enforcement, and justice at scale become infrastructure, and infrastructure embeds values.

Three deployment regimes are emerging. They are simplified here because the mechanism that matters is how deployment gets organized.

China’s approach is state-directed. Build the infrastructure, set the rules, instrument the system, and use the technology for internal stability and external influence. It is deeply attractive to developing economies because it offers a path to modernization without building democratic capacity first. China offers functioning on a faster timeline, even if what you build is something the West would not want to live under.

Europe’s approach is regulate-first. The premise is defensible — systems that touch citizens at scale should be constrained. The EU AI Act formalizes that premise, and in practice, compliance becomes a considerable entry cost. That protects citizens. It may also protect incumbents because the firms best able to pay the compliance tax tend to be those already at scale. Europe’s posture is now inseparable from geopolitics. American companies are increasingly treated as strategic dependencies, and “sovereign cloud” is the policy response — reinforcing the many-entity market, because the compliance regime becomes the boundary the workflow has to live within.

The U.S. approach is innovate-then-regulate. Tolerate volatility in exchange for speed, capture upside early, patch the downsides later with standards and enforcement. That pattern has held for decades of software. It breaks when systems scale faster than oversight.

Policy makers have a harder job than they’re given credit for. In fast channels — media, publishing, advertising — AI produces harm quickly. In slow channels — healthcare, justice, critical infrastructure — the work is gated by institutions, and the integration burden creates a natural barrier. Slow channels are where deployment earns its legitimacy, because the workflow has to survive

to stay alive.

This is why regulation can’t be a complete answer. A policy update is not a deployment strategy. Technology will iterate through hundreds of changes while a committee writes a definition. Even when policy lands, enforcement has to be resourced. The CCP built that resourcing through decades of party construction. New technical opportunities don’t become a contested market first; they become state capability by default.

Deployment regimes compete on outcomes. Which one makes services work at scale without degrading the society running them is not a question any regime has answered yet.



## 12 On Administration and AI

*“When you reduce counting, you increase life.”*

Most writing about AI lives above the machinery. I am less interested in AI as an idea than in where it contacts and replaces work.

Work has transactions. Transactions pass through workflows. Workflows have a counting layer. The counting layer expands volume, but eventually steals capacity from the system it was built to support. That is what administration is when it's healthy, and what bureaucracy is when it metastasizes.

Much of our public discourse about AI skips directly to the extremes. Either the utopia where everything is free, or the dystopia where humans are obsolete. Both are seductive narrative indulgences. The truth is, the world you and I live in is a ledger problem.

We have more people than we used to. We have more digital surfaces per person than we used to. We have more transactions to negotiate because the routing substrate is always on and well optimized. These are good things in the broader context of a society that can share ideas and experiences with one another.

Sophisticated compute is one of the few tools in decades that can carry part of the counting layer without adding humans in proportion.

When the counting layer shrinks, two things happen.

First, capacity appears. People get time back from reconstruction and routing and can spend it on what the work claimed it valued — creativity and collaboration.

Second, service expands. Workflows that were rationed by labor cost or labor scarcity become available to more people. If done earlier, we may see benefits accrued like preventive health.

My optimistic case for AI is boring: make average life better through repeatable improvements. Fewer forms, fewer backfills, fewer handoffs, fewer overworked experts, more throughput, more access.

There are risks.

Output volatility will persist and be damaging where strong demos, even products, get marketed as deployment-ready systems. Displacement will follow as institutions taper hiring and contract, either through more efficient delivery or by flailing in a competitive services market. That displacement hits humans across income, identity, and comfort at the same time.

None of those risks are solved by hoping the technology won't land. The timing for these rewrites is a market question, but the need is structural, and the pressure is compounding.

To get there, it will require measurement, lane models, and the willingness to treat governance as a system with workflows.

Each person will need to find their own explanation of these changes and what it means for themselves and their family.

At Laurel, the mission started as and continues to be an effort to return time. Ryan and I intentionally left the phrase non-prescriptive. People spend time in ways that are subjective and deeply personal. The one thing I feel comfortable asserting is that time spent on administrative reconstruction is rarely the best a human can do, and almost never the best a system should ask for.

## 13 Parting Thoughts

The origins of this essay lie in dozens of conversations, collated into a central thesis about AI's potential to help humanity and the constraints that shape whether integration works.

Had Ryan not invited me on the journey that became Laurel, many of these concepts would have stayed far more abstract. The trenches of building startups create a filter that reveals how frail a smart-sounding hypothesis can be. Or put another way, trust scientists who have broken beakers.

I hope this essay gives you an orientation. Find a workflow you care about. Identify its transactions. Define the counting layer. Ask what you would change if the counting layer shrank by a hundred, or a thousand. Sort out who benefits, who gets displaced, and what you would do about it if you were responsible.

Charlie Munger once told a class of graduating lawyers to deliver to the world what you would buy if you were on the other end. If you are a passionate expert (whether you believe it or not) in your industry, and you have a decade or two ahead of you, the people on the other end of your workflows deserve that. Rebuild from the inside — or take that domain expertise outside and build from the outside with an insider's knowledge.

The administration that once felt restraining can become the fulcrum for your lever.

To the many people in my life, I love you. And in particular, to those who spent time and energy improving this essay, listed in alphabetical order. Ryan Alshak,

Sanchit Bareja, Nicholas Carter, Jonathan Chin, Iza Cottle, Bárbara Dantas, Ashutosh Desai, Natasha Desai, Perla Gámez, Mary Griffith, Mike Griffith, Matthew Joanou, Mitch Katzer, Charlie Melvoin, Patrick J. Nilan, Jacob Sills, Ariane L. Smith, Mike Tobias, Nat Welch, and Eric Zaarour.

# 14 Appendix

## 14.0 CUTTING ROOM FLOOR

Lines that didn't earn a section but kept showing up in conversation. They're here because someone will need them.

*"Procurement is where innovation goes to learn manners."*

*"The counting layer breaks first."*

*"If you can't count the work, you can't modernize the work."*

*"If the workflow is invisible, reform is theater."*

*"The future has more transactions than humans."*

## 14.1 GLOSSARY

The terms I use throughout this essay. This is not a dictionary. It's my operating meaning — what I'm pointing at when I use a phrase, and how I want the reader to interpret it.

## 14.2 CORE UNITS

### Transaction

My meaning: A countable unit of work that must be routed into an outcome the institution recognizes.

Examples: A timesheet line, an invoice, a claim, a permit application, a ticket, a payment, a case update, a compliance attestation.

Related: Workflow, Routing, Counting layer.

### Workflow

My meaning: Routing made enduring. The repeatable path that turns messy reality into something the institution can accept.

The test: If it can't be run repeatedly at scale, it isn't a workflow — it's a one-off.

Related: Transaction, Institution, Defensible.

### Routing

My meaning: The act of moving something to the next place it must go — human, system, ledger, queue, approval, payment, storage.

Related: Workflow, Lane model.

### Digital surfaces

My meaning: The interfaces and tools through which modern work is produced and recorded.

Examples: Email, calendar, docs, chat, tickets, CRMs, payment rails, identity systems, internal tools.

Claim: *More surfaces → more transactions → more routing → more administration unless we redesign.*

Related: Transaction volume, Counting layer.

## 14.3 ACTORS AND SCOPE

### Government

My meaning: A public entity that runs workflows citizens can't meaningfully switch away from, and that can compel participation.

Related: Institution, Gates, Procurement.

### Institution

My meaning: An entity whose workflows have hardened into standards — defensibility, compliance, habit, and inertia.

Related: Gates.

## 14.4 LAYERS OF WORK

### Ledger

My meaning: The artifact the institution treats as "what happened."

Claim: *The ledger isn't recordkeeping; it is how the institution prices and defends reality.*

Related: Counting layer, Administrative reconstruction.

### "At scale, the ledger becomes the work."

My meaning: When transaction volume rises, the effort to produce acceptable artifacts can rival or exceed the effort of the underlying work itself.

### Counting layer

My meaning: The work required to make work legible to the institution so the institution can route, price, and justify it.

Examples: Timesheet reconstruction, categorization, "billable narratives," audit artifacts, compliance checklists, backfills.

Claim: *The counting layer exists because scale exists. It consumes capacity while pretending to be overhead.*

### Administrative reconstruction

My meaning: A specific form of counting layer work: reconstructing a week (or a case, or a project) after the fact so it fits the ledger.

Examples: "Sunday night timesheets," email archaeology, calendar backfills.

### Administration

My meaning: The scaling layer that replaces trust when trust can't scale. It keeps workflows coherent by routing, verifying, recording, constraining, auditing, and paying for work.

Claim: *Administration is necessary, and then it becomes heavy. The line between the two is often volume.*

### Bureaucracy

My meaning: A subset of administration. Administration turned self-protective.

What it looks like: Slow to update, procedure over outcomes, CYA as operating principle.

Claim: *Bureaucracy is not "nothing." It is an evolved response to risk — but it can become a drag strong enough to cap output.*

### Coordination gravity

My meaning: The force that pulls entities toward more administration as transaction volume and internal complexity rise.

Claim: *It's the reason big organizations reach for layers, artifacts, and ritualized coordination.*

Related: Counting layer, Many-entity market.

## 14.5 ADOPTION, RISK, AND DEPLOYMENT TERMS

### Outcome

My meaning: The durable result the buyer purchases (or renews): realization, continuity, compliance, reliability, measurable performance.

Claim: *Institutions buy outcomes, not novelty.*

### Integration

My meaning: The work required to insert capability into a real workflow without breaking defensibility.

Claim: *Capability is cheap. Integration is expensive.*

### Reliability layer

My meaning: The set of mechanisms that make volatile capability repeatable in production.

Examples: Evaluation, monitoring, routing, fallbacks, refusals, auditability, containment.

Claim: *In serious workflows, reliability becomes the work.*

### Downside

My meaning: The scope of harm when a system is wrong, and who carries it.

Use when: I'm talking about risk posture and "what happens when this fails."

Related: Gates, Safety, Human lane.

### Buyer team

My meaning: The combined purchasing decision-makers inside an institution — typically procurement, risk, and IT (plus legal/security, depending on context).

Use when: I want to avoid muddy "buyers/purchasers/procurement" drift.

### Gates

My meaning: The controls an institution puts around workflow change: diligence, audits, security review, procurement logic, compliance posture, and continuity requirements.

*Claim: Gates are often rational. They become pathological when they stop filtering for quality and start filtering for change. Related: Rake, Procurement.*

### Rake

My meaning: A structural posture that keeps workflows embedded and competition weak — framed as "safety" or "standardization," monetized as suite purchasing and vendor lock-in.

*Claim: When gates become profitable, the institution stops modernizing its workflow and starts selling its inertia.*

### Vendor capture / Regulatory capture

My meaning: When a single vendor becomes the default path for modernization, and "standards" become whatever that vendor ships.

*Claim: It can feel like progress and still set up the next generation of ossification.*

### Lane model

My meaning: How automation enters real work:

Automated lane: High-volume transactions handled with minimal human involvement.

Hybrid lane: Compute assists; humans stay in the loop for confirmation, review, escalation.

Human lane: Exceptions, contested judgment, and outcomes where the cost of being wrong is unacceptable.

*Claim: The lanes exist early; what changes is the proportion of transactions that can safely move into each lane.*

### Reversibility

My meaning: How easily a workflow step can be unwound once shipped.

*Claim: Reversibility determines how aggressively you can widen automation without turning one mistake into a cascade.*

## 14.6 AI AND COMPUTE TERMS

### Artificial Intelligence / AI

My meaning: A bucket term. Useful socially, imprecise technically.

### Sophisticated compute

My meaning: Machine learning capability deployed into workflows at a meaningful scale — available as infrastructure, not rare R&D.

Why I use it: "AI" has become shorthand for everything. "Sophisticated compute" keeps the discussion grounded in deployment and capacity.

### Machine learning

My meaning: The technical family that produces these capabilities (including LLMs), without implying science fiction.

### LLMs

My meaning: Language models that provide flexible input handling and generative output.

*My stance: Powerful, visible, volatile — useful, but not a universal deployment primitive.*

### Diffusion-family systems

My meaning: Generative systems that produce media at scale (images, video, etc.).

## 14.7 THE TRIADS I REUSE

### Routable / Repeatable / Defensible

Routable: The output can move through the next step of the workflow without manual reinterpretation.

Repeatable: The outcome holds under volume; it's not a one-off success.

Defensible: The institution can justify the output to clients, auditors, regulators, or courts.

*Claim: Mature workflows don't run on "good answers." They run on these three properties.*

### Revenue / Trust / Compliance

Revenue: The workflow connects to money (pricing, realization, recovery, budget).

Trust: The workflow connects to confidence in outcomes (clients, citizens, internal stakeholders).

Compliance: The workflow connects to rules and enforcement (audits, security, oversight).

*Claim: These are the arteries. If you touch them, you inherit their risk profile.*

### Latency / Accuracy / Safety

Latency: The cost and speed of producing an output (compute, time-to-answer, throughput).

Accuracy: The degree to which outputs match the intended truth or requirement in context.

Safety: The degree to which the system can be deployed without unacceptable harm (to users, institutions, or society).

*Claim: Real systems choose tradeoffs. "All three" is a sales line, not an engineering plan.*

### Income / Identity / Comfort

Income: The ability to support yourself and your family.

Identity: The narrative of worth attached to work.

Comfort: The stability of routine, status, and predictable rules.

*Claim: This transition hits all three at once, which is why people defend work they often don't even like.*

## 14.8 MACRO ARCS

### Digitization → Globalization → Centralization

My meaning:

Digitization: Work artifacts and workflows moved onto computers.

Globalization: Work routed across borders and time zones because the substrate became digital.

Centralization: The winners of that era concentrated data, distribution, and power by owning the routing layer.

*Claim: That arc grew GDP and capability, and pulled value upward. AI changes the advantage structure again.*

## 14.9 ANALOGIES

### Shipping container

My meaning: Standards are ecosystems. You don't change one component without changing what depends on it.

### Propeller

My meaning: Better geometry becomes possible when manufacturing changes; new forms require the rest of the system to certify and support them.