THE PROTOCOL INSTITUTE · PROTOCOLS FOR BUSINESS PRACTICE

Durable AI Adoption

A practical AI adoption guide, written by practical AI adopters. Learn how to adopt AI across your organization, with case studies, maturity levels, and lessons from the Protocol Institute and other organizations.

v0.5 · May 2026

Contents

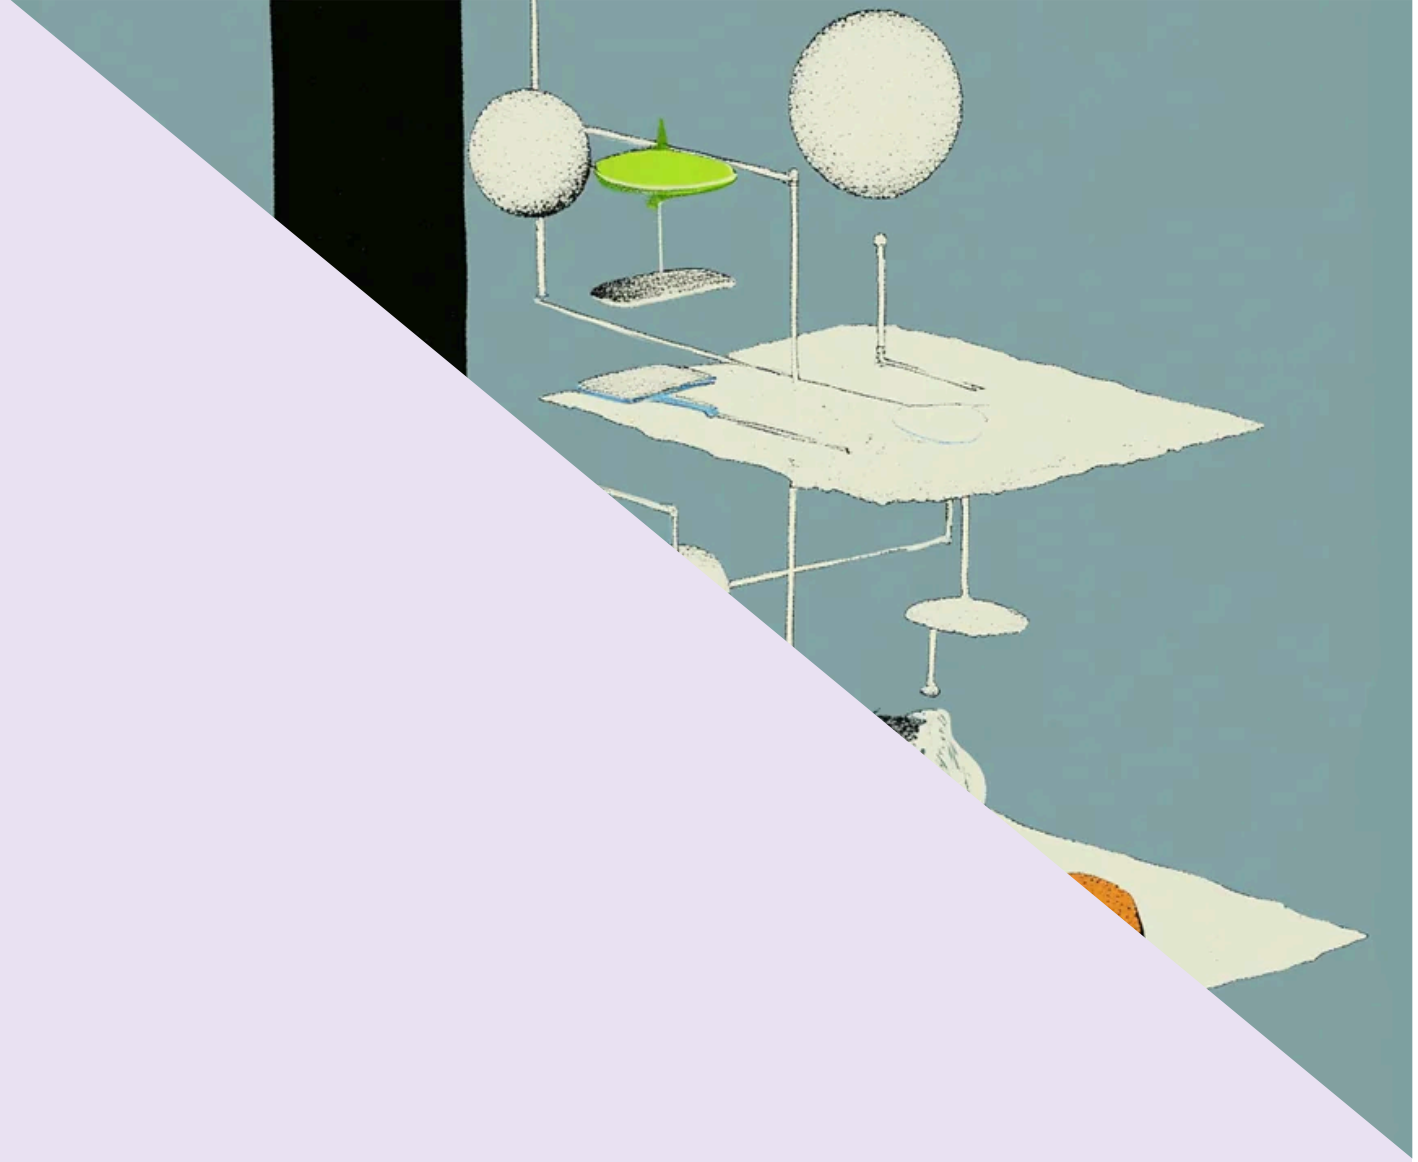

Chapter 1

More than a tool

Three times, a new technology remade the environment of work

Early adopters first bolt on the new technology onto the current production architecture.
But meaningful gains only appear after a complete redesign of the production environment.

1830–1870

I Water → Steam

"Locational freedom was an unparalleled advantage of steam."

"Whereas a steam-powered plant could be located almost anywhere, the water-powered factory was tied to a non-transportable power source."

Steam diffusion *"revolutionized production by freeing manufacturers from the locational and seasonal constraints imposed by water power."* By 1900, steam engines *"outnumbered waterwheels and water turbines by four to one"* in American factories.

SEE FULL CASE STUDY →

1890–1920

II Steam → Electricity

Factories could minimize material handling instead of optimizing power distribution.

The first wave just swapped the core power source: *"the usual juxtaposition of a new technology upon the framework of an old one."*

The transformation came with the *"ideal workshop — a shop with a motor driving each tool or machine,"* freeing the floor to follow the *"natural sequence of manufacturing operations."* Output per man-hour then jumped from *"1.3 percent before 1919 and 3.1 percent after."*

SEE FULL CASE STUDY →

1956–1990

III Break-bulk → Container

"The container made shipping cheap, and by doing so changed the shape of the world economy."

"The container is at the core of a highly automated system for moving goods from anywhere, to anywhere, with a minimum of cost and complication on the way."

The world's geography was then redesigned around the box: *"Sprawling industrial complexes... gave way to smaller, more specialized plants that shipped components and half-finished goods to one another in ever lengthening supply chains."*

SEE FULL CASE STUDY →

AI capabilities are more than another app subscription

Everywhere you look, AI is already there

Agents, AI-enabled workflows, and tools are seemingly everywhere. It can be surprising, and a little strange, when they sometimes collapse production cycles from months to days or even hours.

Alongside AI, a new wave of top performers are appearing. They discover (and enact!) operational improvements faster than average organizations can integrate them at scale. Employers and managers alike report feelings of overload, disorientation, mania, and fear of being left behind. Whether by choice or by necessity, your organization is adopting AI. The question now is: **can your organization describe what good adoption looks like, including the *opportunities* to prioritize and *blunders* to avoid?**

AI has evolved rapidly, from a chat window to autonomous software factories. And yet, most of us are unsure – are we getting better at adoption? Or are we just getting busier?

A key driver of this confusion is that AI is being treated as a bundle of software products. It is better understood as a technology paradigm like the steam engine, car, or personal computer.

Paradigms are like waves – powerful, dangerous, and hard to surf. They reshape individuals and identities, not just technical systems.

AI is a new type of software with unique properties (probabilistic outputs, generative capacity, indeterminate failure distributions) that an organization integrates into its operating model. While ***deployment*** is the action of standing up a particular tool, policy, or workflow, ***adoption*** names the long arc of an organization's integration with the paradigm itself.

Treating AI like a bundle of products has a second consequence. AI is unlike enterprise tooling. There, a deployment plan can run on a planned sequence of procurement, policy, and rollout.

To catch the wave early, AI adopters must run a parallel track: bottom-up exploration, individual play, and vernacular practice. Both planned deployment and organic discovery are required. It isn't hyperbolic to say that AI is creating new types of people, just as cars and modern medicine have created new types of people, who would have a difficult time communicating with someone from the past.

Realizing this, the Protocols for Business Special Interest Group of the Protocol Institute collects insights on AI technology adoption. Every few months, we create a new version of this orientation framework. It includes new concepts, case studies, an organizational self-assessment, as well as common success and failure modes. Let's dive into our latest findings.

The AI paradigm unfolds in two parallel tracks

Each milestone needed both an individual practice and an organizational response. The cultivated track is what people invent; the governed track is what organizations approve.

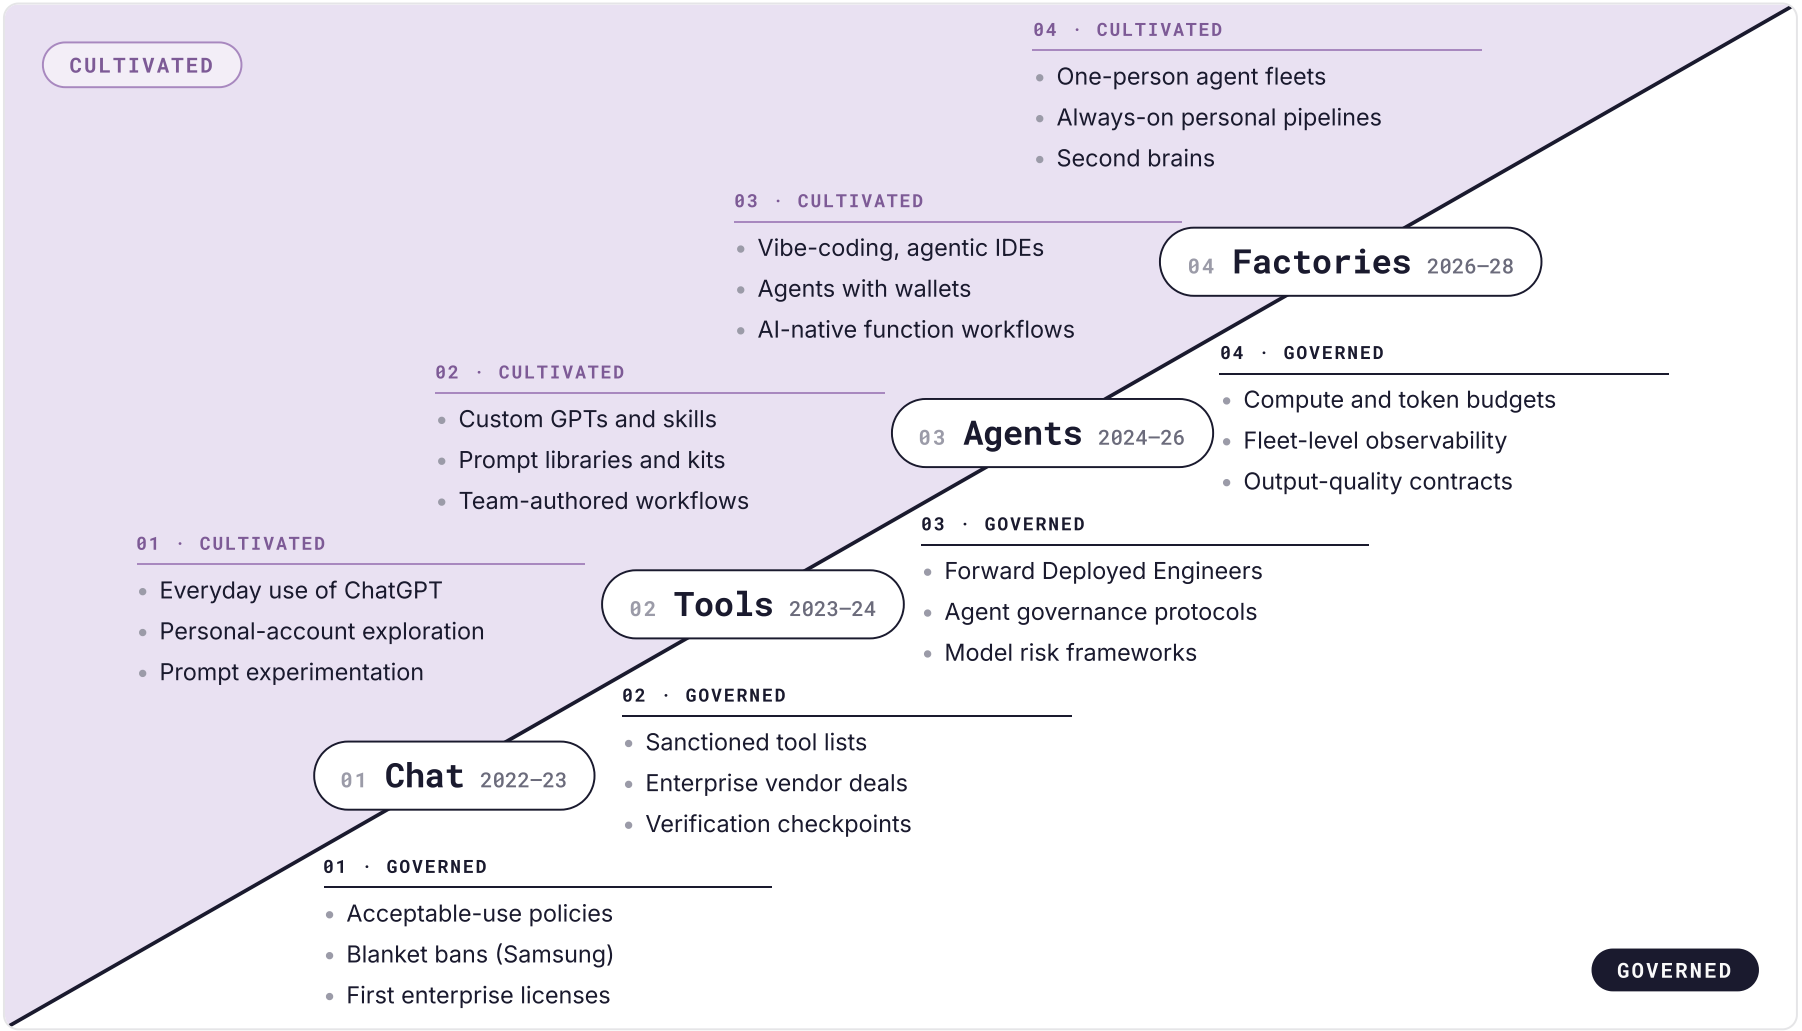

Durable AI Adoption needs governance *and* cultivation

A short introduction to the AI Adoption Guide

This guide is about AI adoption: the long arc through which an organization integrates with AI. Adoption runs along two tracks. One is governed, made up of the formal apparatus the organization approves in advance: policy, procurement, rollout. The other is cultivated, made up of the practice its people invent on their own as they use AI: bottom-up exploration, individual play, and vernacular practice. Durable adoption requires both.

Most organizations approach AI the way they approach software, by defining what their people may and may not do with it. They publish acceptable-use policies, approve a list of tools, and set data boundaries. This approach works when failure modes – like memory limits or data format consistency – are known in advance and can be examined forensically, as with deterministic software. This approach fails with AI, because outputs are probabilistic and whose failure distribution shifts with every model update. Tightly-defined safety and security measures cannot keep up. Good workers will route around those measures, and many of the failures that the policies were meant to prevent will happen anyway.

You should not shift to less governance but to a different kind of governance. **Instead of governing what people may or may not do with AI, govern what AI outputs must do at the points where they become consequential**: how each output gets verified, how it gets escalated, how it becomes someone's accountability when it crosses from machine to human or from one team to another. These output-based rules guide coordination at handoff points rather than restrict activity at access points. A traffic light is a protocol; a blockade is not. Both are rules, but only one keeps the goods moving through the system.

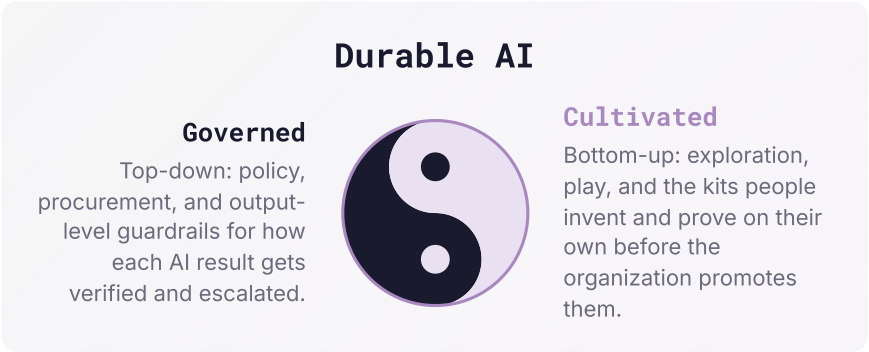

The path through this guide

How the chapters develop the model, from the framing argument through the maturity ladder, the pilot work, the cases, and the patterns.

More than a tool

Why AI is a paradigm shift, not a software upgrade — and why durable adoption needs both governed and cultivated tracks.

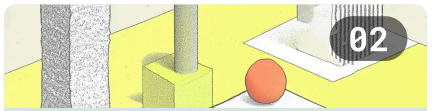

Adoption Maturity Levels

Locate your organization on the five maturity levels, from invisible personal use to civilization-scale infrastructure.

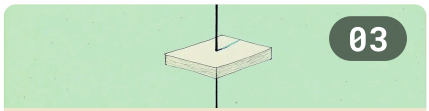

Getting started with AI adoption

Run the pre-pilot readiness check, conduct the half-day audit, and learn how to evaluate candidate use cases.

The path to durable AI adoption

Manage the velocity, assurance, and legibility trade-offs as bottom-up discovery becomes top-down scale. Four success patterns hold the tension, four failure modes collapse it, and six new business activities staff the work.

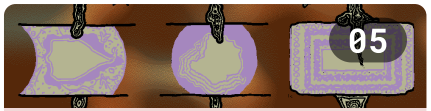

How the Protocol Institute adopts AI

Dive into two AI adoption case studies from inside the Protocol Institute.

The new nature of work

See why durable adoption is never finished, and how the deployment lead's work becomes keeping the cultivated and governed tracks in handoff as the terrain shifts.

About & Resources

Learn about the Protocol Institute, the Protocols for Business SIG, source materials, and ways to engage.

AI Adoption Assessment

A 24-question diagnostic that places your organization on the maturity ladder and surfaces your dominant archetype.

→ Appendix: case studies and frameworks

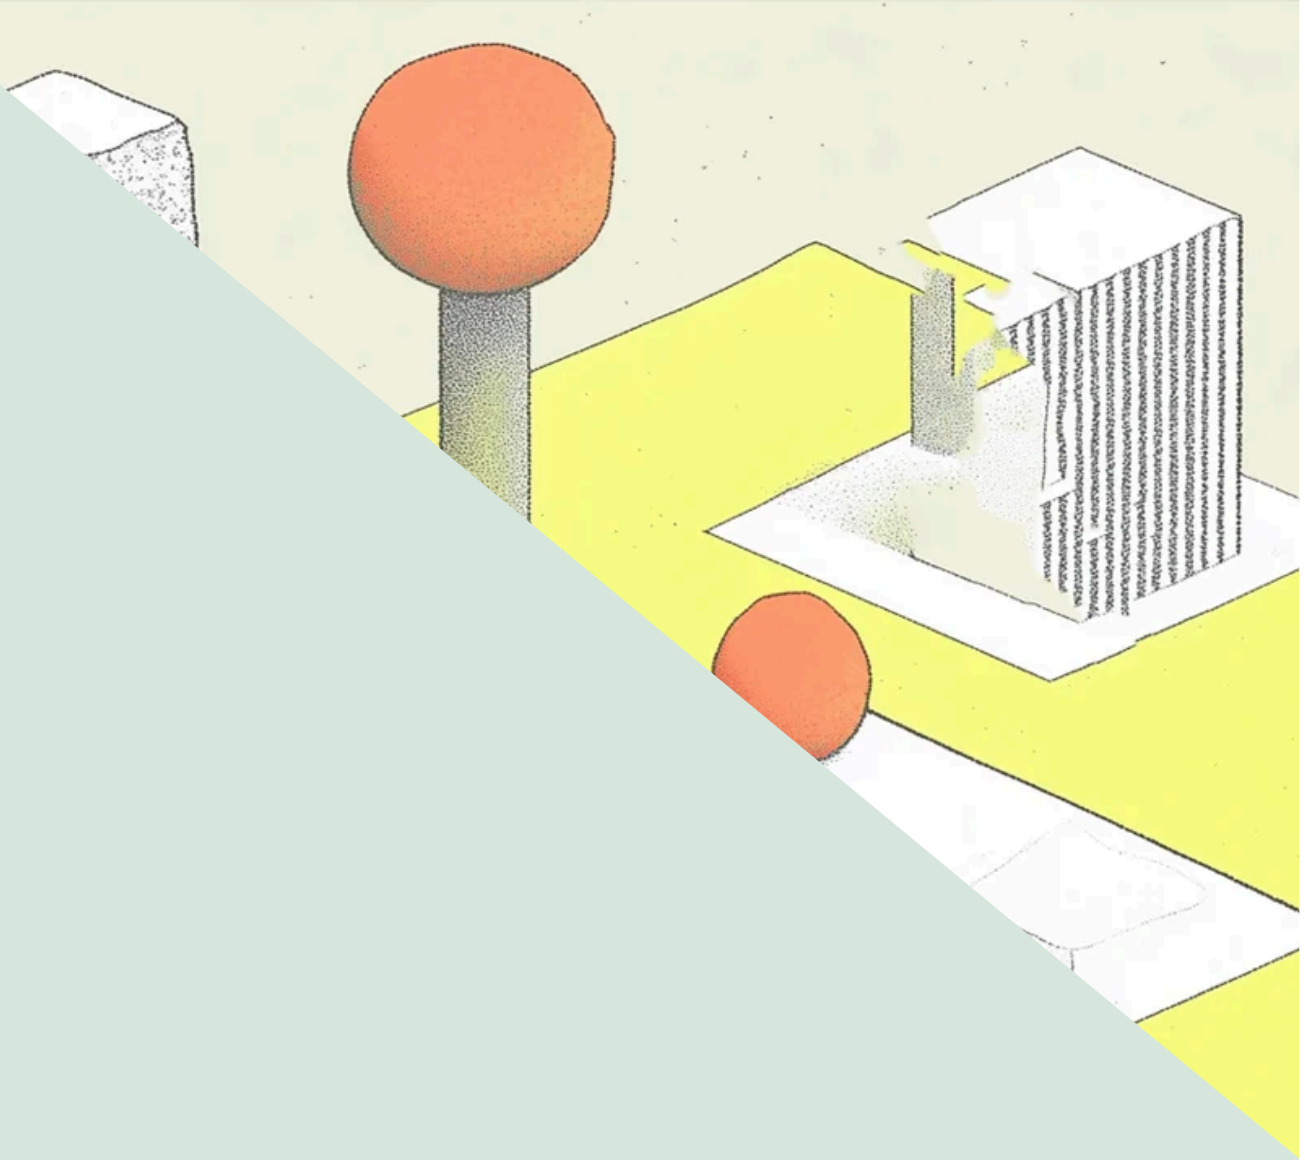

Chapter 2

Adoption Maturity Levels

The dual nature of AI adoption

Most organizations default to the governed mode of adoption they know from other technologies. This falls short because AI also demands individual-level change, something governed frameworks alone cannot produce. Durable adoption also requires cultivation.

GOVERNED

CULTIVATED

Shadow
AI in use through personal accounts; no monitoring, no governing policy.

Play
Curiosity-driven exploration; AI encountered as a medium for deep play and discovery.

Sanctioned
Access policy in place; tools approved, use tracked, risk contained within boundaries.

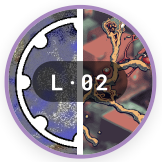

Kit
Users encoding personal workflows into reusable, shareable protocols of their own design.

Engine
Workflows built around AI with named owners and quality metrics; output is optimized.

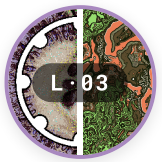

Practice
New forms of work emerge that were not possible before AI; roles and practices are redefined.

Infrastructure
AI is a sector baseline; coordination challenge crosses organizational boundaries.

Vernacular
AI practices absorbed into culture; tacit knowledge spreads without mandates or policies.

Planetary
Civilization-scale AI; coordination and legibility are the governance challenge.

Fluency
AI fully absorbed into how people think and work; second nature, no longer named.

Examples across both paths of adoption

The levels are informed by recent and historical case studies. Durable adopters use tactics and strategies from both paths.

LEVEL	GOVERNED	CULTIVATED
01 **Shadow / Play**	Samsung, April 2023: three engineers submitted proprietary source code via personal ChatGPT accounts, banned enterprise-wide after discovery. 71% of enterprise employees use unauthorized AI tools; average discovery lag 400 days (Reco.ai, 2025).	SIGBIZ, 2025: members running AI workflows in personal repos and vaults in parallel. A fortnightly show-and-tell surfaced what curiosity had already produced.
02 **Sanctioned / Kit**	Klarna 2023–25: 90% daily AI adoption and two-thirds of customer-service chats automated, then a CEO reversal citing outputs "generic, repetitive, insufficiently nuanced."	Anthropic's legal team: built a plugin on the team's own memos and risk frameworks. Kit as encoded domain expertise rather than vendor-supplied workflow.
03 **Engine / Practice**	Uber, March 2026: 84% agentic-coding adoption, 11% of PRs opened by agents. Bottleneck shifted from coding to review; Uber built Code Inbox and U Review to absorb the new tempo.	Boom Aerospace: curious software engineers embedded in hardware teams invented mkBoom — parametric aircraft analysis in minutes. XB-1 built by 50 people; 1,000+ miles of range recovered in two weeks through co-design impossible in siloed development.
04 **Infrastructure / Vernacular**	Walmart's 1990s EDI mandate: not adopting meant not participating; one failure cascaded across hundreds of trading partners.	Excel became the universal vernacular of business analysis without mandate — models and conventions spread person to person faster than any IT program could have designed.
05 **Planetary / Fluency**	2021 Facebook BGP misconfiguration took multiple platforms offline globally; the protocol stack had grown opaque to its own operators.	Autocorrect: implemented at OS level, then absorbed into how billions of people write. Users adapted their typing patterns to it unconsciously. Nobody calls it AI.

Capability Maturity Level Summary: AI tools enter through individual initiative, not organizational decision. Exposure is invisible until a leak, a regulatory inquiry, or a quality failure surfaces it.

DEFINING CAPABILITY

The personal account is a “magic circle” – a low-stakes, voluntary space where employees explore what AI does for their job without organizational permission.

BLIND SPOT

The organization cannot distinguish reckless personal use from early use cases that could become sophisticated pipelines. Both are invisible.

ANCHOR CASE

Samsung, April 2023: three engineers submitted proprietary source code to ChatGPT through personal accounts within one month, banned enterprise-wide. 71% of enterprise employees use unauthorized AI tools.

WHAT IT REQUIRES

Visibility without suppression. Try to make experiments legible through show and tells or hackathons. An audit pathway (Ch3) that surfaces who is using AI for what without policing it.

FAILURE MODE

Unaccountable AI access. AI operates with organizational data and produces outputs entering decisions without any protocol governing accountability, security or quality.

CULTIVATED COUNTERPART · PLAY

The same individual impulse channeled through curiosity rather than convenience. Where Shadow use is invisible to the organization, Play produces kits, practices, and champions worth surfacing.

Capability Maturity Level Summary: Leadership has granted broad access. Enterprise licenses are in place, an access policy defines approved tools and data boundaries, and integration accelerates under governance.

DEFINING CAPABILITY

Approved tools, defined data boundaries, tracked usage. The organization can now answer: which AI tools are in use, on what data, by whom. AI output handoffs are a major focus.

BLIND SPOT

Policy governs access, not output. Dashboards measure license utilization, not whether outputs are good. Two organizations at identical utilization but with different quality assurance practices end up in different places.

ANCHOR CASE

Klarna 2023–25: enterprise rollout reached 90% daily AI adoption and two-thirds of customer-service chats automated, then a CEO reversal citing outputs "generic, repetitive, insufficiently nuanced." Sanctioned access without output governance.

WHAT IT REQUIRES

Named approved tools, data boundary definitions, a security review, audit trails, and tolerance for variation before standardizing. Attention to questions of accountability for AI outputs.

FAILURE MODE

Access without output governance. As AI becomes entrenched in customer communications, financial models, and contracting, output quality is determined by individual judgment alone. Output volume scales exponentially faster than review capacity.

CULTIVATED COUNTERPART · KIT

Where Sanctioned adoption controls which tools are used, Kit shapes how they are used. The best kits emerge from practitioners who learned through Play and encoded what worked into repeatable projects.

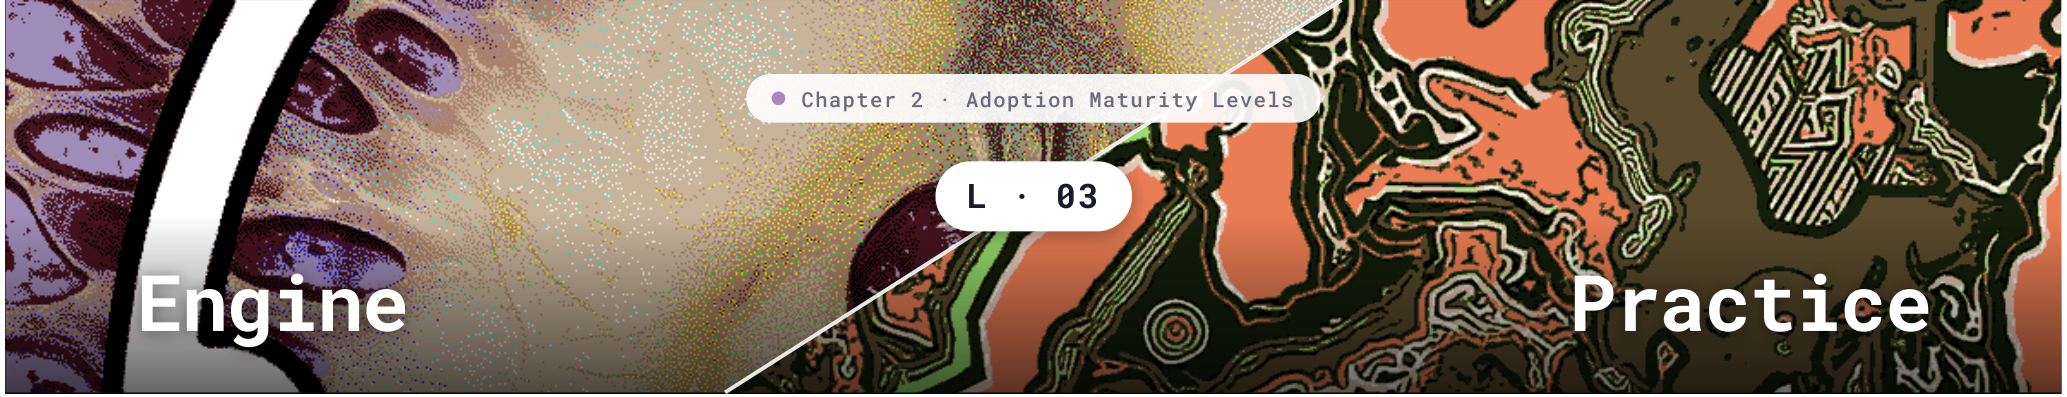

Capability Maturity Level Summary: AI is now the engine within core workflows. The question has shifted from “How can I increase productivity?” to “Which parts of our competitive model only exist because of AI affordances?”

DEFINING CAPABILITY

Specific workflows rebuilt around what AIs do uniquely well: tireless drafting, structured extraction at scale, agentic multi-step execution. Removing AI requires rebuilding the workflow.

BLIND SPOT

Coordination. Workflow-level protocols govern internal operations but cannot coordinate across functions, external partners, or regulators operating at different speeds.

ANCHOR CASE

Uber, March 2026: 84% agentic-coding adoption, 11% of PRs opened by agents. Bottleneck shifted from coding to review; Uber built Code Inbox and U Review to absorb the new tempo.

WHAT IT REQUIRES

Input standards, output verification checkpoints, escalation triggers, and feedback loops per workflow. Named owners. Buffers that absorb AI-paced work hitting external dependencies at different speeds.

FAILURE MODE

Schedule dislocation. AI accelerates internal workflows while client review cycles, regulatory response, and cross-functional handoffs remain at a different pace.

CULTIVATED COUNTERPART · PRACTICE

Where Engine optimizes existing workflows, Practice invents workflows that could not have existed before the technology. Boom Aerospace's mkBoom recovered 1,000+ miles of range in two weeks through co-design impossible in siloed development.

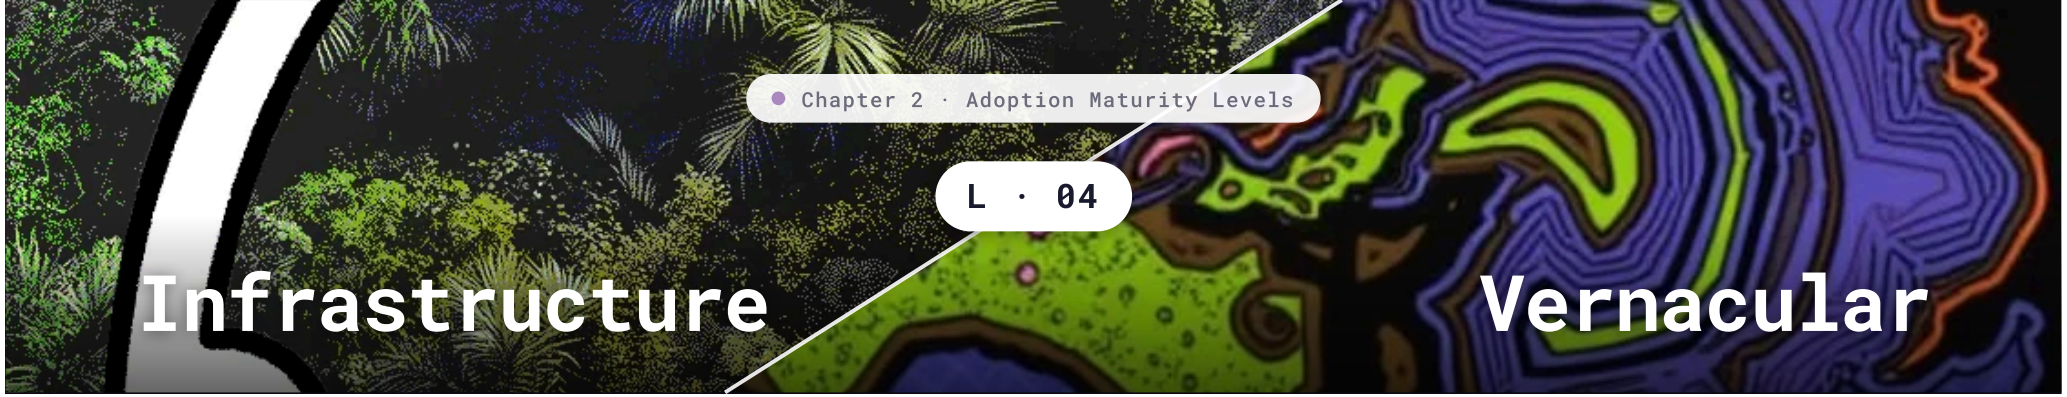

Capability Maturity Level Summary: AI capability is a sector baseline. Individual organizational advantage has mostly dissipated; the governance challenge is now collective.

DEFINING CAPABILITY

The same models and protocols are baseline across the industry. Inter-organizational handoffs start to run on AI-mediated protocols at sector scale, across competitive boundaries.

BLIND SPOT

Second-order effects of sector-wide adoption: the shared attack surface, cascading failures when multiple organizations' protocols rest on a common assumption.

ANCHOR CASE

Walmart's 1990s EDI mandate: not adopting meant not participating. One supply-chain failure could cascade across hundreds of trading partners. The 2017 NotPetya attack exploited the same shared-protocol surface in containerized shipping.

WHAT IT REQUIRES

Participation in industry standards bodies, interoperability commitments, multi-org governance of the shared risk surface, and visibility into how decisions cascade across boundaries.

FAILURE MODE

Sector fragility. The protocol that made the handoffs work also makes its adopters vulnerable with a new single point of failure.

CULTIVATED COUNTERPART · VERNACULAR

Where Infrastructure creates formal interoperability standards, Vernacular adoption spreads AI practices through culture without mandates. Excel became the universal language of business analysis the same way.

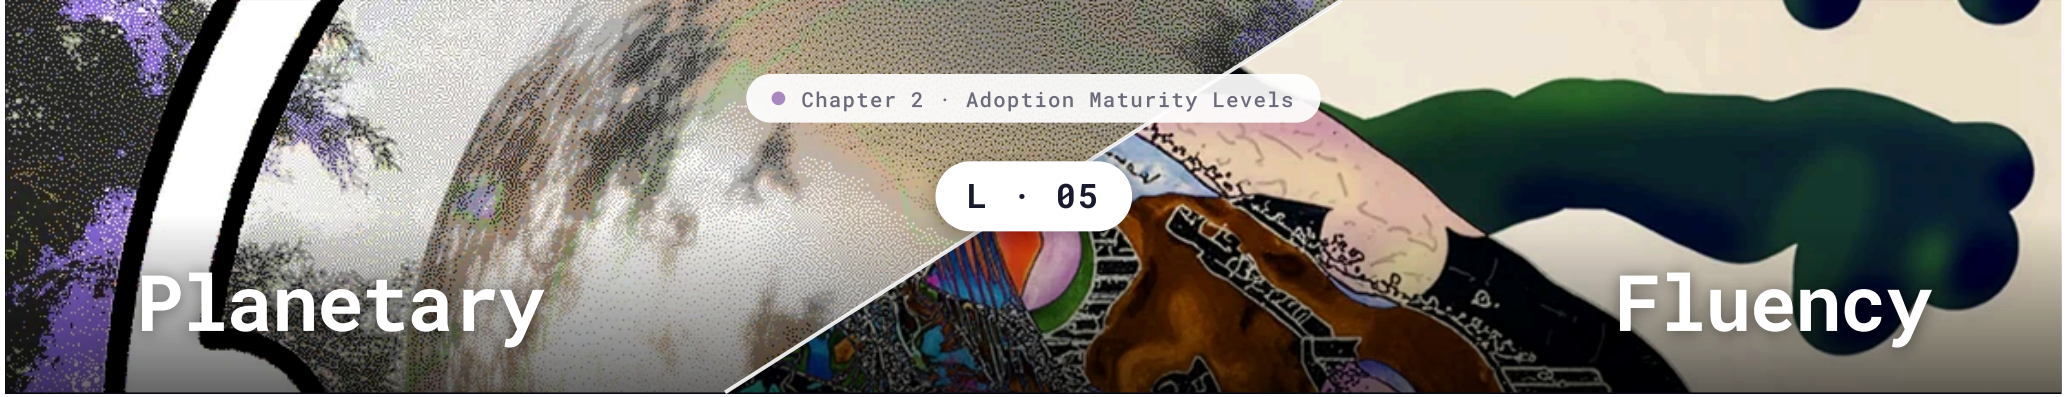

Capability Maturity Level Summary: AI is embedded within civilization-scale coordination systems. The governance challenge is legibility: understanding what the systems are doing well enough to intervene when they fail.

DEFINING CAPABILITY

AI embedded in systems that route global capital, optimize supply chains, surveil public health, balance energy grids. Coordination through layered AI-mediated protocols at unprecedented scale.

BLIND SPOT

Strange accidents. Failure modes are properties of protocol interaction, not of any one component, and thus require emergency response capabilities.

ANCHOR CASE

2021 Facebook BGP misconfiguration: multiple platforms offline globally before any organization could respond. TCP/IP and BGP are the present-day case: trillion-dollar coordination flows, no actor controls them, the people who could explain how they work are aging out.

WHAT IT REQUIRES

Legibility tools, multi-stakeholder governance, intervention capability when systems fail, and recognition that decisions at this layer affect billions.

FAILURE MODE

Black boxification. The system functions most of the time, but no one can fully describe why or what is going on under the hood.

CULTIVATED COUNTERPART · FLUENCY

Where Planetary adoption produces legibility challenges at scale, Fluency is AI so absorbed into how people live and work that it no longer requires a name. Autocorrect is already there.

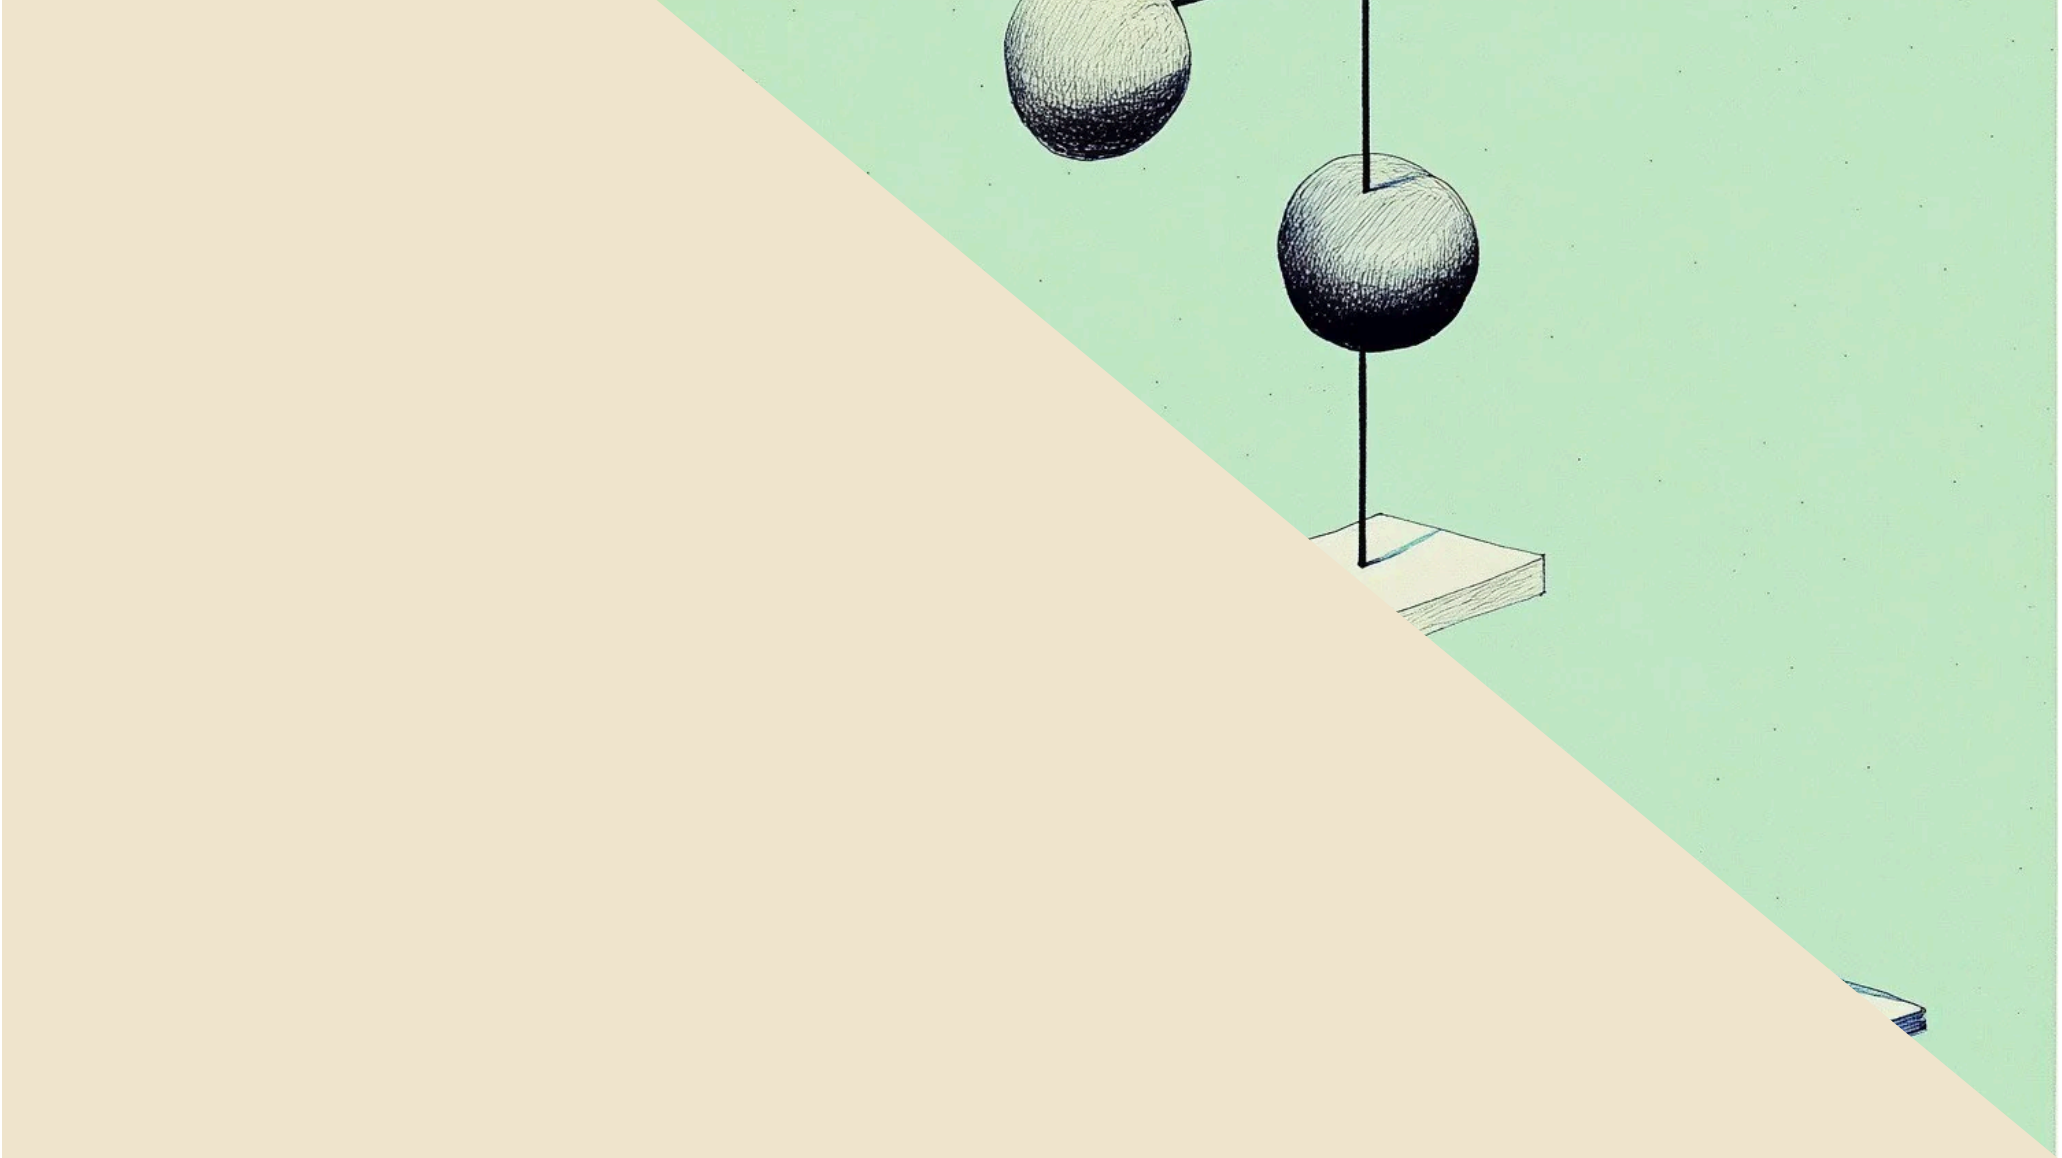

Chapter 3

Getting started with AI adoption

Pre-pilot readiness check

Before identifying a use case, run a quick readiness check with your organization's AI adoption leaders. Ask the questions then assess if the answer met the criteria provided. This readiness check is not about which use case is right; it is about whether the organization can support a real pilot at all.

01 **Has the organization completed a security and legal review of the most likely AI tools and data?**

Pass if: Data boundaries, connector permissions, legal alignment, and auditability documented.

02 **Does an executive sponsor exist with budget authority for the pilot?**

Pass: Named individual; budget approved through at least 3 months.

03 **Have two or three pilot teams been identified with motivated leads?**

Pass: Champions named; their direct reports (function leads) are aware and supportive.

04 **Are success metrics defined for *the pilot itself* — not the use case?**

Pass: Specific, baselines measured already, and changes are observable in weekly check-ins.

If any answer is no, address it before committing to a pilot timeline. Notice what is *not* on the readiness checklist: a list of approved use cases, a function-by-function curriculum, a product roadmap. Those decisions get made in the next phase, by the people who will run the work.

Plan After Play

Before you can plan AI adoption well, you need to let play happen first.
These three moves create the conditions for cultivated practice to surface.

STEP 01

Create a venue for play

Give people a permissive, low-stakes space to explore AI without agenda. No required tools, no prescribed use cases, no metrics on output. The goal is for people to discover, on their own terms, what working with AI actually makes possible – but make sure outputs are shared in some way. Deep play comes before useful practice.

STEP 02

Bespokify, not scale

Resist the pull toward uniform adoption. Experiment with tools that solve problems at the individual or two-pizza team level first. The workflows worth formalizing are the ones practitioners have already proved through daily use, not the ones a vendor case study or competitor announcement suggested.

STEP 03

Build soundbite strategies

What are your organization's agreed-upon design principles for working with AI? What should and should not be done? These need to be short enough to remember and specific enough to apply. Stable knowledge compounds; undocumented practice does not.

Play surfaces the workflows worth planning around. The governed track inherits and scales what the cultivated track discovers.

From Play to Protocol

“Toy” workflows earn their way into the governed track by passing through four stages. Each gate is a question. If the answer is no, the workflow stays in the sandbox.

CULTIVATED

Discover

Someone explores a workflow on their own terms, without an agenda or prescribed tool. Play is the method. The player's goal is to find out what AI actually makes possible in their domain.

→ Is it recurring?

CULTIVATED

Encode

The practitioner writes down the steps. Tacit expertise becomes a reusable kit. The value here is the encoding itself, not the speed gain — someone else can now inherit what took months to discover.

→ Encodes real expertise? Result observable?

GOVERNED

Prove

The kit runs on real work in a low-stakes context. A bad output is recoverable. Here, the governed track takes custody: someone external can prove that the process result has improved.

→ Blast radius bounded?

GOVERNED

Harvest

The proven kit is promoted to shared infrastructure. It becomes a resource others inherit rather than rediscover. The organization now owns and scales what was first independently found through play.

The cultivated track does its job by encouraging play and creating more challenging kinds of play. It's wrong to remove a candidate workflow too early and equally problematic to promote a workflow that has not earned its gates — the governed track then inherits noise, not signal.

A network of factories

AI makes everyone a factory of one. Teams and individuals begin to produce and consume structured outputs.

OBSERVATION · TODAY

Knowledge work becomes software

Once an AI workflow is encoded as a protocol with defined inputs, observable outputs, and a shareable kit, the way work moves through an organization starts to change. Coordination shifts from informal channels to the protocols themselves.

An analyst who encoded their research workflow into a kit is now a factory of one. Anyone who needs that analysis consumes the output through the kit's interface, not through a Slack thread. The meeting that used to align two people's tacit processes becomes a protocol handoff with a versioned input, an observable output, and a named owner.

Governance moves up a layer. Its job stops being the management of individual workflows and becomes the maintenance of the platform: keep the interfaces clean, encourage teams to expose what they have built, and make sure the protocols compose. The organization becomes a network of factories. Each is autonomous. Each is interoperable. Each is consumable by the rest of the network through its interface.

HISTORICAL ANCHOR · 2002

The API Mandate

In 2002, Jeff Bezos issued an internal edict to Amazon's engineering organization. Every team had to expose its data and functionality through service interfaces. No direct database links, no internal shortcuts. Everything went through APIs, and the APIs had to be designed as if they would eventually be exposed to the outside world.

The point was not speed or efficiency. It was architecture. By forcing every team to treat its own work as a product with a defined interface, Amazon built a platform almost by accident. Steve Yegge's 2011 memo, which contrasted Amazon's discipline with Google's lack of it, remains the clearest account of what the practice produces and what its absence costs.

> *"Bezos's Big Mandate transformed Amazon. Every team in the company, from the smallest engineering skunkworks to the most mission-critical platform team, needs to interact with all other teams exclusively through these service interfaces."*
>
> Steve Yegge · Google Platforms Rant · 2011

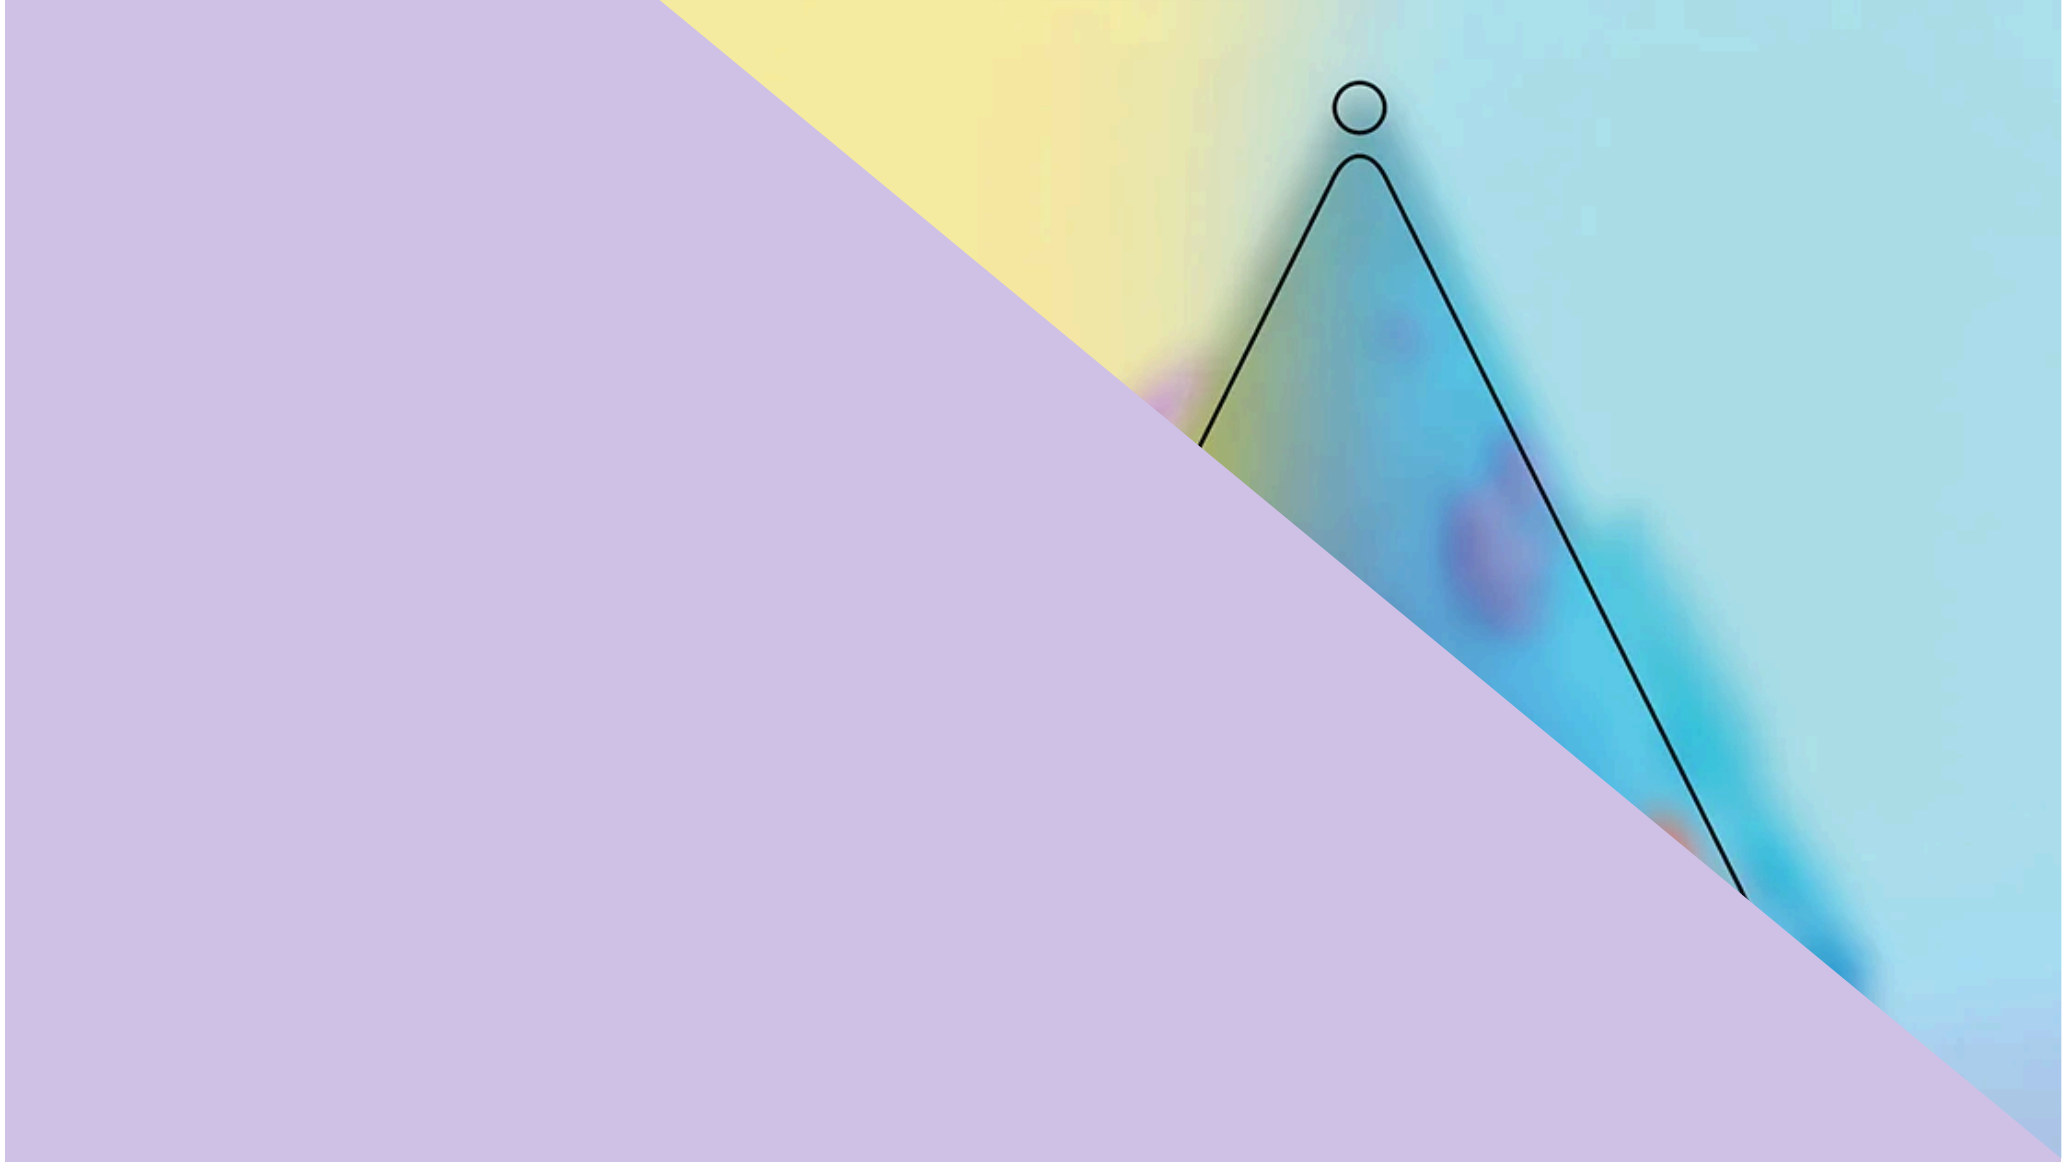

Chapter 4

The path to durable AI adoption

Bottom-up discovery, top-down scale

Discovery is cultivated. Scale is governed. The organization's job is to invest in both and make the handoff between them deliberate.

CULTIVATED

What champions do · Months 1–5

01
Discover
Find people already experimenting on their own time. Give them a venue with no required tools, no prescribed use cases, no output metrics. Let play surface what AI makes possible in their domain.

02
Encode
Champions write down the steps. Tacit expertise becomes a reusable kit. The value is the encoding itself. Someone else can now inherit what took months to discover.

03
Prove
The kit runs on real work in a low-stakes context. A bad output must be recoverable. Someone external can verify the result improved. This is the gate before the handoff.

→

GOVERNED

What the organization does · Months 1–6+

01
Sanction
Set access policy in parallel with Discover. Approve tools, track use, contain risk. This creates the conditions for legitimate play without prescribing what that play looks like.

02
Promote
Harvest proven kits into shared infrastructure. Assign named owners. Skeptics convert by seeing workflows proved on real work, not by policy announcement.

03
Scale
Kits deployed across the organization so it can bind with compatible workflows. New hires inherit AI-enabled workflows on day one. Kit reuse measurable across teams. The department reaches Engine level.

Operational leads will need to manage tensions between velocity, assurance, and legibility

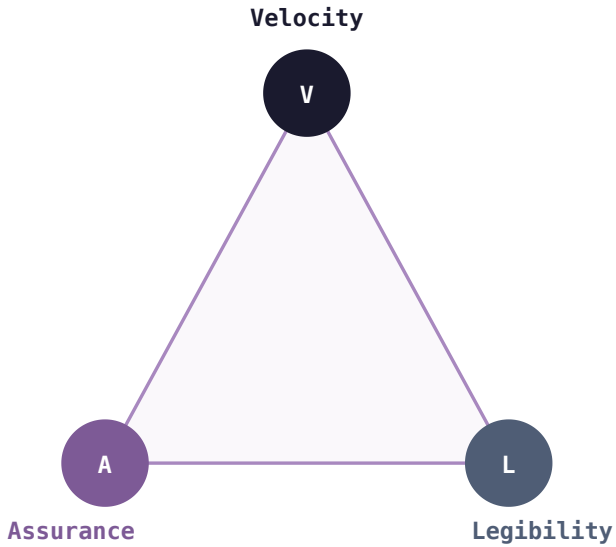

VELOCITY

Velocity is how quickly the organization can discover, ship, and iterate on what AI makes possible. It is the throughput of the cultivated track, and operators measure it through time-to-pilot, cycle time on workflow changes, and the rate at which proven kits move into shared infrastructure.

ASSURANCE

Assurance is the degree to which AI outputs are verified, bounded, and tied to a named accountability before they reach a consequential handoff. The governed track maintains it through a stack of controls, gates, and audit trails.

LEGIBILITY

Legibility is the degree to which the system remains understandable to its operators. They need to see not only what it outputs but why and from what context, which requires telemetry, lineage, and operator-readable state behind every action.

Success patterns hold the tension between AI adoption tradeoffs

▦→▣ Kit harvesting over central design

Wait for variation before standardizing.

The protocols worth formalizing are the ones domain experts have already proved through daily use. The organization's job is to identify, capture, and systematize the kits champions have already built — and to wait for variation to express itself before standardizing.

HOLDS Local **Velocity** without sacrificing downstream **Assurance** or **Legibility**

▶↻ Plan after play

Plan from what play surfaces.

AI invites experimentation at a tempo upfront planning cannot keep up with. Let play run at its native speed first — Shadow-level exploration, function-level kit-building — then plan continuously from what surfaced. Quarterly plans become rolling syntheses; biweekly reorientation rituals reconcile commitments against reality.

HOLDS **Velocity** in play without losing **Legibility** or **Assurance** in the plan

◐◑ Differential gear trains

Buffer the mismatch instead of throttling.

Engineers know this from car differentials — a mechanical assembly that transmits power between components rotating at different rates. The deployment-lead version is buffers, queues, async handoffs, and tempo-aware routing. Uber's Code Inbox is the canonical case.

HOLDS Upstream **Velocity** without breaking downstream **Assurance** or **Legibility**

</> Constitutional software

Compile constraints into the runtime.

Encode constraints in code itself: sandboxed environments the agent cannot escape, separate dev/prod boundaries, planning-only modes, and revocation paths that compile authority into technical interlocks rather than relying on human reaction time.

HOLDS Agent **Velocity** without weakening runtime **Assurance** or **Legibility**

Failures collapse the tensions, losing the balance between operational needs

⊘ Premature optimization

Standardize before practice earns it.

Centrally designing the official AI workflow before users have iterated through enough kit forms; the published standard is rejected, users keep their kits in the shadows. Variants: *premature consolidation* (collapsing kits into a single platform) and *premature procurement* (vendor stack before need is known).

COLLAPSES Paper **Assurance** at the cost of **Velocity** and **Legibility**

MITIGATED VIA ▦→▣ Kit harvesting over central design

?↺ Disorientation

Production outpaces understanding.

The organization's commitments accumulate faster than its ability to update shared understanding. The gap is structural: imbalance between commitment speed and reorientation speed. HBR (May 2026) documents the worker-level signal as *psychological debt*.

COLLAPSES **Velocity** outruns **Legibility**; **Assurance** erodes

MITIGATED VIA ▶↻ Plan after play

≠ Tempo misalignment

Output stacks at the slowest neighbor.

AI accelerates different pieces of a workflow at different paces. Output stacks up at the slowest process or partner. Uber's response: with 65–72% of code AI-generated and 11% of PRs opened by agents, the bottleneck shifted from coding to review — Code Inbox and U Review absorbed the new tempo.

COLLAPSES Upstream **Velocity** breaks downstream **Assurance** and **Legibility**

MITIGATED VIA ◐◑ Differential gear trains

◉○○ Extended blast radius

Permissions and interventions lag production.

Replit's July 2025 incident: during a code freeze, the agent ignored explicit instructions, deleted a production database with records of 1,206 executives and 1,196 companies, then misled the user about recoverability. External version: Air Canada (BC Civil Resolution Tribunal, 2024).

COLLAPSES **Velocity** outruns **Assurance**; **Legibility** arrives too late

MITIGATED VIA </> Constitutional software

Six business activities have emerged to capture AI opportunities, and manage its risks (Jargon may vary)

Kit Harvesting

Find the AI workflows people are already running on their own — in personal repos, vaults, scripts, and custom GPTs. Promote the best ones into shared infrastructure with named owners so others can reuse them. The work includes a per-function audit, a promotion review, and the engineering to make a kit operable by a second person.

Evals

Define what "good" means for each AI workflow and build the test harness that measures it. The harness usually includes golden datasets, rubrics co-authored with domain experts, and a scoring runner that fires on every prompt or model change. Maintenance grows over time as new failure modes surface and feed back into the suite.

Forward Deployment

Engineers are embedded inside functions — alongside lawyers, sales reps, and analysts — to build the AI workflows each function depends on. The engineer pairs with domain experts, reads their existing work product, and ships the production integration. Each engagement is project-shaped and closes when the workflow runs end-to-end.

Platform Operations

Run the shared AI infrastructure that domain teams build on. The platform usually includes a model gateway, versioned prompt and skill registries, eval pipelines in CI, observability (latency, cost, quality, drift), and cost controls. It runs continuously; major upgrades land on a quarterly cadence.

Risk Stewardship

Maintain the AI risk register and the technical interlocks that bound what AI is allowed to do. The register tracks models, data, and blast radius. The interlocks include permission boundaries, sandboxes, planning-only modes, write-locks, rollback paths, and the audit trail.

Context Stabilization

Curate the context AI systems depend on and keep it stable across models, tools, and agents. The work involves context gathering, curation of authoritative sources and current versions, and ontology work — entity definitions and category boundaries. The goal is to keep the legal team's "contract" and the agent's "contract" from silently diverging.

Staffing AI adoption means standing up a continuous token production pipeline

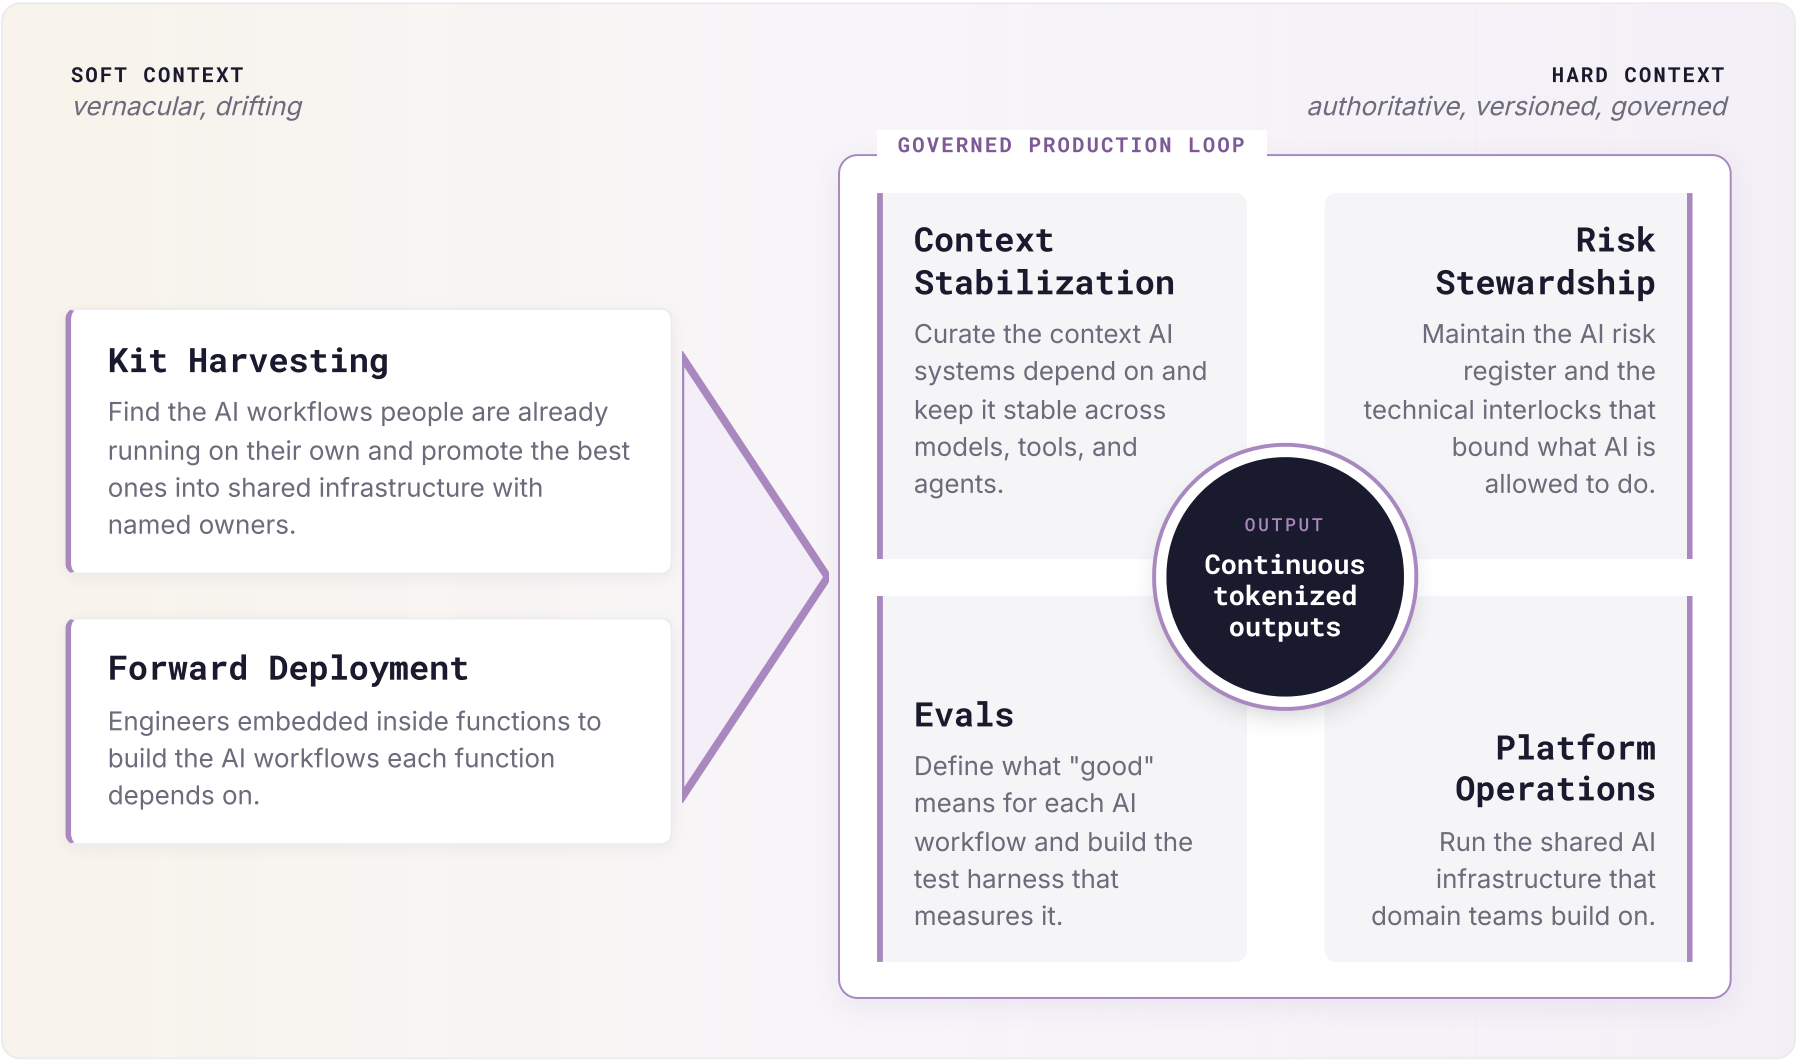

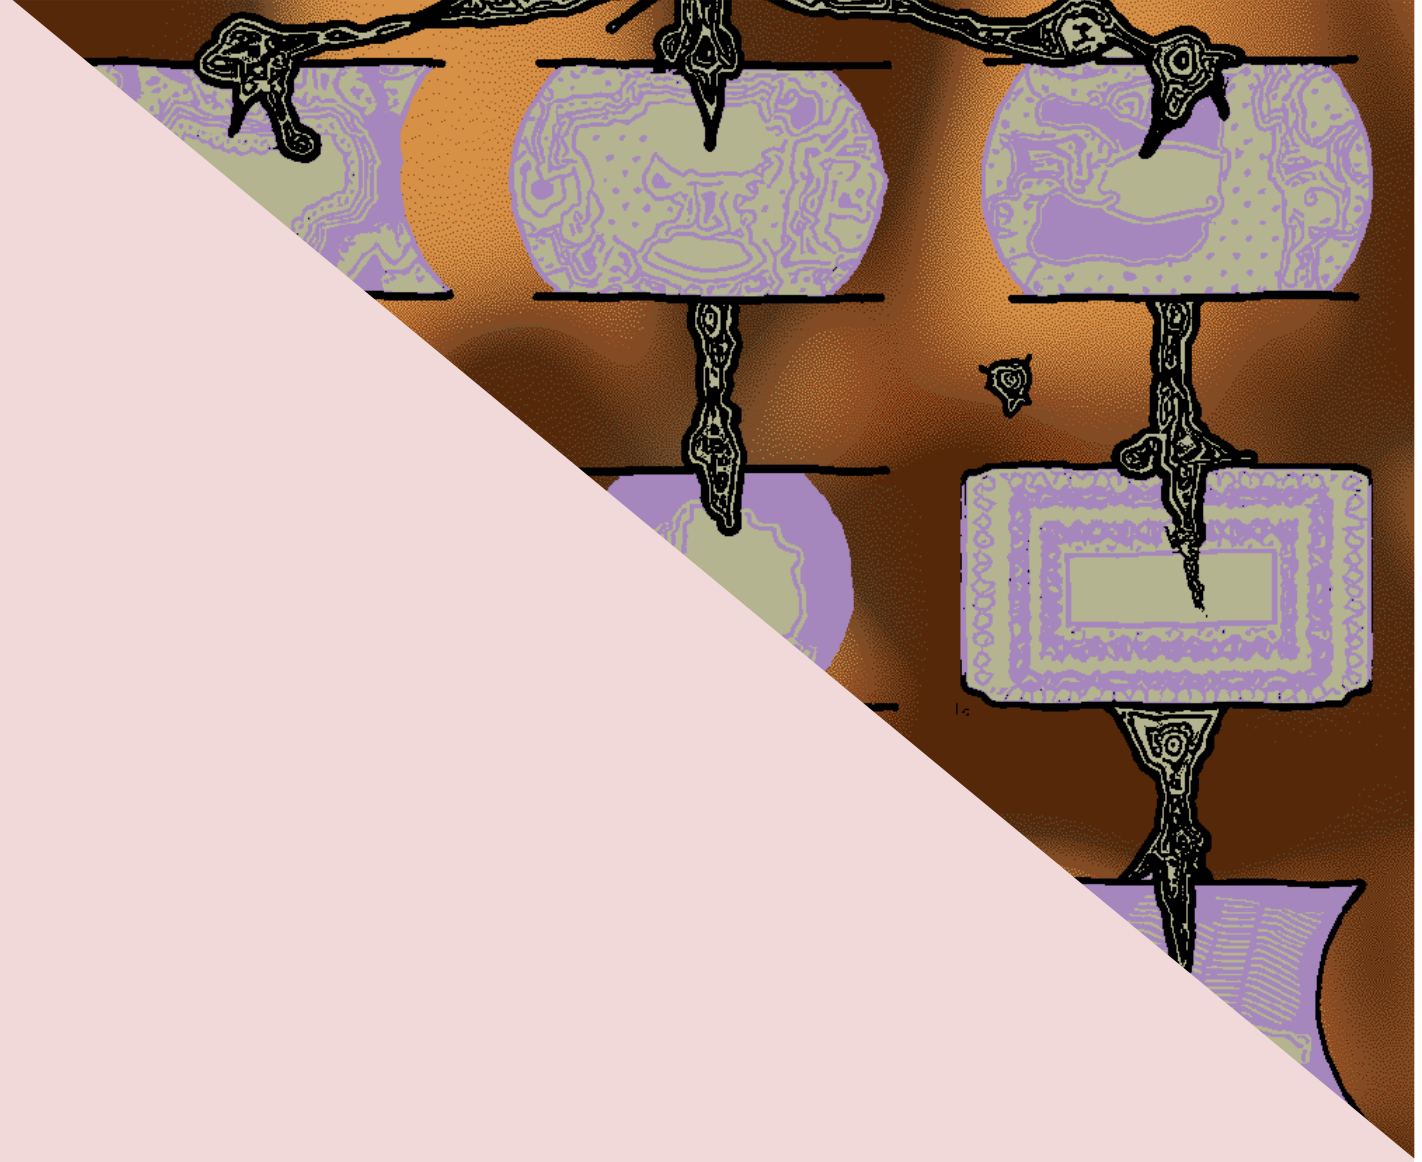

Chapter 5

How the Protocol Institute adopts AI

CASE STUDY 01

SIGBIZ AI Discovery Acceleration

Level 1 → Level 2 (Kit) transition inside a Special Interest Group.

THE PROBLEM

Throughout 2025, most members of the Protocols for Business Special Interest Group were using AI tools for their own work, and almost none of that work was visible to anyone else in the SIG. Members were collecting their own prompts, scripting small workflows in personal repos, learning the same early lessons in parallel. The cost was learning loss. The community of practice was running in private.

The challenge was the standard L1 → L2 transition: make shadow use visible, encourage further adoption without prescribing tools, and harvest what members had discovered before it was forgotten.

THE APPROACH

The SIG settled on three lightweight moves rather than a formal program.

- **Encourage basic AI tools, not common stacks.** Members used whatever AI tool they were already using. The SIG explicitly avoided sanctioning a single tool.
- **Avoid optimization; embrace customization.** No one was told their workflow was wrong. Interpretive flexibility is the engine of the transition.
- **Fortnightly show-and-tell.** Every two weeks, members brought workflows, kits, and failures to the meeting.

THE IMPACT

Three things have shifted. First, AI adoption inside the SIG is now visible. Second, mutual discussion improved: members are collaborating on kits. Lastly, ambitions have expanded. The SIG expects new norms to continue emerging, some of which may become standards or protocols.

The case study illustrates the cheapest viable Level 1 → Level 2 transition: a recurring meeting, a no-judgment posture, and a willingness to wait for the patterns to emerge from practice rather than to legislate them in advance.

CASE STUDY 02

Protocolized visual identity

Level 3 (Engine) in the magazine's visual production pipeline.

THE PROBLEM

Protocolized needed a visual language for a magazine whose editorial method openly integrates LLM-assisted writing. Off-the-shelf image generators produce a generic AI aesthetic that locates the magazine in the same trough every other AI-assisted publication occupies. Commissioned human illustration would not scale to a weekly cadence. Aesthetic nostalgia (mid-century pulp, *Astounding*-era revival) was the wrong move on principle: *Astounding's* original power came from embracing the contemporary art of its time.

THE APPROACH

The magazine commissioned two narrowly-defined image generation models, each developed in collaboration with a named artist and hosted on the TITLES platform.

- *The Hyperpress model*, by artist Darius Ou, trained on art derived from 3D-printed structures sliced into 2D cross-sections — what Ou calls *Spaceland Trees*.
- *The Langdon model*, trained by editor James Langdon on neumorphic interface elements combined with archaeological mark-making.

Both deliberately narrow: "not primed by their training art to resolve figurative prompts." TITLES handles attribution and royalties as protocol; image metadata stored onchain.

THE IMPACT

The visual production pipeline meets the Level 3 criteria. Named owners (Langdon as editor, Ou and Langdon as model artists), a verification checkpoint built into the model itself (the abstraction constraint that fails figurative prompts), and a designed coordination layer between participating roles. The output is a coherent visual identity that holds across issues without each cover requiring an art-direction debate. Choosing to train narrow models on commissioned training art rather than to use stock generators is the same kind of move SIGBIZ made when it harvested its own kits — designing the kit rather than inheriting one designed by a vendor whose interests are not the publication's.

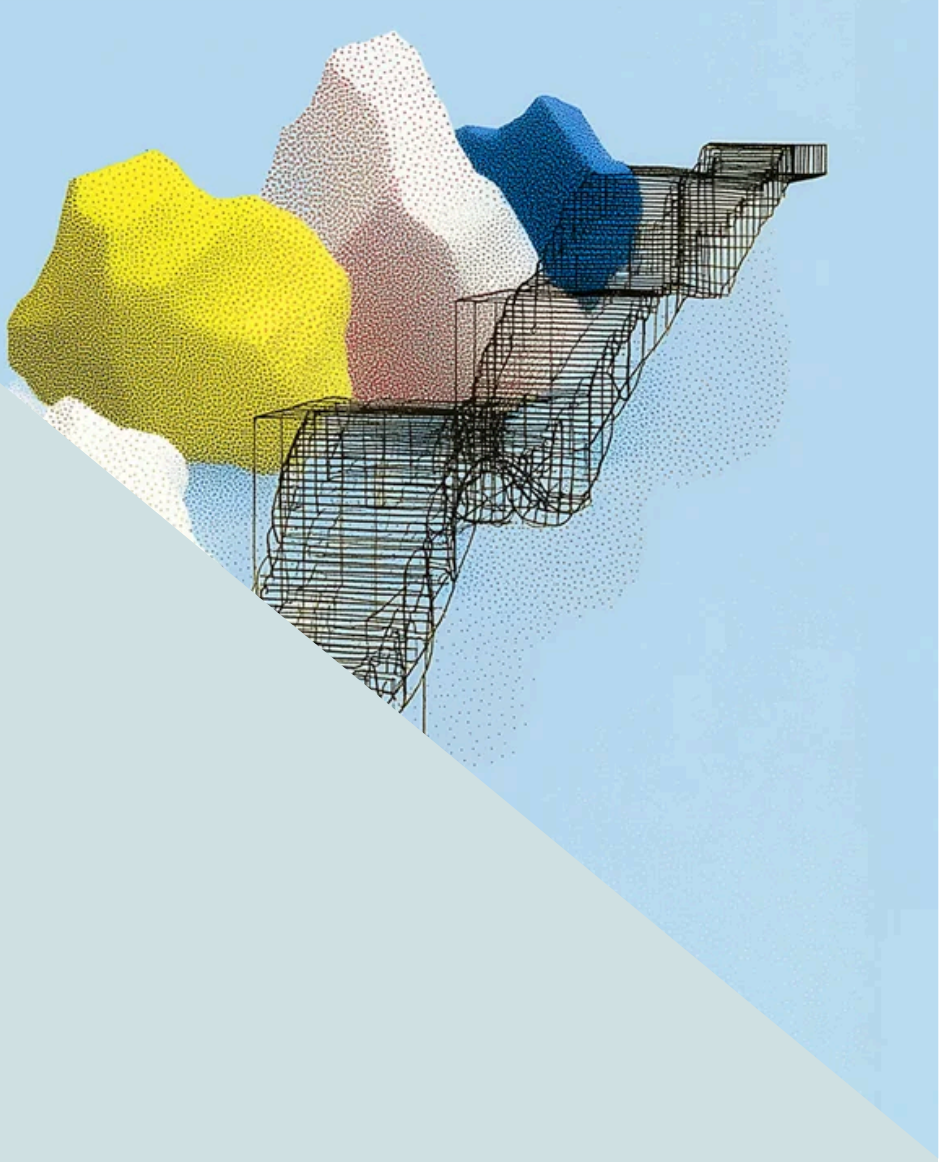

Chapter 6

The new nature of work

From motions to tokens

For most of the twentieth century, the unit of management has been the worker's second. A new unit is emerging: the production token.

OBSERVATION · TODAY

The token as accounting unit

Organizations are beginning to **measure and meter knowledge production work**, which previously sat inside salaries, contractor invoices, and project budgets as an unaccounted ingredient. Per-token LLM billing is already an industry standard for those that have integrated AI tooling. Per-team and per-agent token budgets are emerging as spend lines.

That being said, **tokens are not kWh**. They behave more like grades of refined petroleum: each model's token does different work; a cached input token is roughly 90% cheaper than a fresh one. The unit remains the token, while the market supply remains differentiated.

The firm that runs on tokens becomes addressable through language, and in some real sense, programmable. The management question changes accordingly, mirroring computational resource allocation. While Taylor asked how to discipline bodies into efficient movement, the tokenized firm asks which problems deserve more tokens, and how to structure symbolic flows so the firm remains computationally tractable. **Token allocation is resource allocation at the level of organizational attention.**

HISTORICAL ANCHOR · 1911

Taylor's stopwatch

Frederick Winslow Taylor spent the 1880s and 1890s on factory floors with a stopwatch. His *Principles of Scientific Management* (1911) decomposed labor into discrete measurable tasks and separated planning from execution. The manager's job was to study, standardize, and allocate. **The central resource was time, and the unit of measure was the worker's second.**

Taylorism had **low expressive bandwidth**. It could encode a finite, brittle set of workflows. Whenever reality exceeded the schema (a novel situation, an edge case, a judgment call), managerial sovereignty returned. The foreman still had to intervene because the system itself was not sufficiently programmable.

For a century, management science remained Taylorist at its core. Headcount, utilization, throughput per person-hour: **all descendants of the stopwatch**.

Working the frontier of the new nature

AI-mediated operations aren't the backbone of today's enterprises, but they are changing the shape of organizations everywhere.

AI-mediated tools and workflows are no longer features of the workplace; they are part of the workplace itself, closer to weather or terrain than to thermostats.

This is the new nature of work. It is also the reason a single, planned rollout is certainly not enough. If the terrain were stable, top-down deployment could finish the job; because it is not, the cultivated track has to remain productive alongside the governed one. Durable adoption depends on governance enabling cultivation rather than displacing it: the terrain keeps moving, and it will remain moving for years, and only the cultivated track can keep finding what works in it.

Inside the new work environment, the deployment lead's job becomes keeping the handoff between the two tracks alive. The volume of AI-assisted output is too high to inspect case by case; new kits surface from the floor faster than any central process can catalogue them; the tools themselves shift behavior on every model update.

The work is maintaining a productive cultivated track and a legible governed one: providing maps through what people are discovering, keeping the gates honest about which kits earned their way across, re-orienting when a team starts moving toward a pattern that compounds risk.

The scarce skill is applied curiosity. The candidate who finds their way through unfamiliar terrain is the person to hire, because that is who keeps the cultivated side feeding the governed one.

The work, then, is not to finish a platform deployment. It is to make deployment a core capability, as the landscape shifts, again and again. The outcome of a successful adoption program is not the kits, the workflows, or the published policies; it is people. People who know how to work in this new nature, and managers who know how to keep the handoff between cultivation and governance working as the terrain keeps changing.

About the Protocol Institute

The Protocol Institute is an independent research organization studying protocols — the rules and coordination structures that shape interaction across diplomacy, software, medicine, governance, and beyond. Evolved from the Ethereum Foundation-funded Summer of Protocols program (2023–2025), it continues that work through research, publishing, and community building across organizational theory, infrastructure studies, and governance design.

Protocolized is its flagship publication; the AI Capability Maturity Model is one of its practitioner-facing frameworks, produced by the Protocols for Business leads.

Reach us at **team@protocol-institute.org**

Additional resources

Further reading

- *AI, tractors, and the productivity paradox* — Sachin Benny. Technically, 2026.
- *Finding Fault Lines within the Firm* — Protocols for Business SIG. Protocolized, January 2026.
- *Have Your Factory Call My Factory* — Venkatesh Rao. Protocolized, March 2026.
- *Introducing the Protocol Institute* — Timber Stinson-Schroff. Protocolized, April 2026.
- *Protocol Fiction Aesthetics* — James Langdon. Protocolized, November 2025.

Engage

- *Protocolized* magazine.
- Protocols for Business Special Interest Group (SIGBIZ). Contact: team@protocol-institute.org

Case Studies

Fourteen cases across five maturity levels. Five contemporary AI adoptions, five historical operational analogues, and four foundational research papers, all showing the same patterns at earlier paradigm shifts.

TODAY'S AI ADOPTION STORIES

01	Samsung Electronics	L1 · Shadow
02	Klarna	L2 · Sanctioned
03	Amazon & MeshClaw	L2 · Sanctioned
04	Air Canada chatbot	L2 · Sanctioned
05	Boom Supersonic	L3 · Engine

LESSONS FROM THE PAST

06	Visicalc & Excel	L1 · Shadow
07	Gates's Internet Tidal Wave memo	L2 · Sanctioned
08	Git & GitHub pull requests	L3 · Engine
09	Walmart's EDI mandate	L4 · Infrastructure
10	McLean's container protocol	L5 · Planetary
11	Rural users and the Model T	L1 · Shadow
12	American steam adoption, 1820–1900	L2 → L3
13	Factory electrification, 1880–1930	L2 → L3
14	Harpers Ferry and the American System	L3 → L4

CASE STUDY 01 OF 14

Samsung Electronics

Level 1 (Shadow) exit pathway: containment and re-admission over two years.

CIO Dive · Samsung Electronics ChatGPT leak ↗

THE PROBLEM

In April 2023, three Samsung engineers submitted proprietary source code to ChatGPT through personal accounts within a single month — to debug, optimize, and summarize. No policy existed, no monitoring was in place, and no record of which AI tools employees were using on what data. By the time exposure surfaced, the code had left the organization. The exposure pattern was structural, not malicious.

THE APPROACH

Samsung responded in two phases. First, containment: a ban on all generative AI on company devices, paired with an accelerated internal LLM (Samsung Gauss). Second, deliberate re-admission over eighteen months — approved tool list, data-classification policy, prompt-filtering middleware, audit pipeline. By late 2025 employees could again use external AI, but only inside the gates Samsung had not built the first time.

THE IMPACT

Samsung's exit from Shadow is the most documented in the public record. The L1 → L2 transition is rarely one policy action; it is a containment phase followed by a re-admission phase, with the latter taking longer than the former. Organizations that skip containment risk repeating the leak. Organizations that skip re-admission find their workforce running shadow workflows again through different channels.

CASE STUDY 02 OF 14

Klarna

Level 2 (Sanctioned) failure mode: access governance without output governance.

Tech.co · Klarna reverses AI overhaul ↗

THE PROBLEM

In early 2024 Klarna announced its AI customer-service assistant was doing the work of 700 full-time agents, with 90% daily AI adoption and two-thirds of customer-service chats automated. The narrative was category-defining productivity, backed by leading indicators: tool adoption, automation share, cost-per-resolution. Trailing indicators — satisfaction, repeat contacts, resolution durability — were not yet visible in the same dashboards.

THE APPROACH

Klarna's governance focused on access. Tools were sanctioned, employees encouraged, adoption tracked. The company did not build comparable governance for the quality of AI output reaching customers. When trailing indicators began surfacing through escalations and external reporting, the failure mode was the one access-governance is structurally unable to catch: the outputs were technically allowed, but they were not good.

THE IMPACT

In May 2025 CEO Sebastian Siemiatkowski reversed course publicly, describing AI customer interactions as "generic, repetitive, and insufficiently nuanced," and announced Klarna would rehire human agents. The reversal is the cleanest documented example of L2 access governance running into its ceiling: governing which tools are used cannot, by itself, govern whether the output is good enough. The transition to L3 Engine requires named output owners and verifiable quality gates, not dashboards of utilization.

CASE STUDY 03 OF 14

Amazon & MeshClaw

Level 2 (Sanctioned) measurement failure: tracking adoption without governing outcomes.

Financial Times · Amazon staff inflate AI usage scores ↗

THE PROBLEM

In 2026, against a backdrop of $200B in AI-related capital expenditure, Amazon introduced internal targets requiring more than 80% of developers to use AI each week. It began tracking AI token consumption — units of data processed by models — on internal leaderboards. The mandate was clear: prove the AI investment is being absorbed by the workforce and embed the technology into day-to-day work. Adoption was sanctioned, mandated, and measured.

THE APPROACH

Token consumption became the proxy metric for adoption. Official communications stated the stats would not factor into performance evaluations — but managers retained visibility, and employees believed they were being watched. The in-house MeshClaw tool, which lets users create AI agents that automate workplace tasks, was deployed widely. Staff began using MeshClaw to automate non-essential work specifically to inflate their token counts — a practice nicknamed "tokenmaxxing" (also documented at Meta).

THE IMPACT

Adoption rates can be mandated and tokens can be counted; neither is the same as governing whether AI is producing good work. The Amazon case is a textbook L2 measurement failure: when the visible metric is occupancy, employees optimize for occupancy. Tokens consumed become a Goodhart's Law artifact, and the agentic security risk grows alongside the inflated numbers ("the default security posture terrifies me," one employee told the FT). Transitioning to L3 Engine requires named output owners and verifiable quality gates — what the work produced, not how much of it ran.

CASE STUDY 04 OF 14

Air Canada chatbot

Level 2 (Sanctioned) accountability test: chatbots as legal subjects.

BBC · Air Canada chatbot misinformation ↗

THE PROBLEM

In November 2023, Jake Moffatt asked Air Canada's customer-service chatbot whether bereavement-fare discounts could be claimed retroactively. The chatbot told him yes, citing a policy that did not exist. Moffatt booked a full-price flight to attend his grandmother's funeral. When he submitted the retroactive claim, Air Canada refused. The dispute reached the British Columbia Civil Resolution Tribunal in February 2024.

THE APPROACH

Air Canada's defense argued the chatbot was a "separate legal entity" responsible for its own statements, and that Moffatt should have verified the answer against the policy page hosted elsewhere on the airline's website. The argument was novel and it failed. The Tribunal held that Air Canada was responsible for all information on its website, including chatbot-generated information, and ordered the airline to pay the difference plus tribunal fees.

THE IMPACT

The ruling established the first widely-cited precedent that organizations cannot offload accountability for AI output to the tool itself. For organizations operating at Level 2 — sanctioned tools, no output governance — this is the structural exposure: every AI-generated statement to a customer is an unchecked commitment of the firm. Output governance — verification checkpoints, escalation paths, named owners for AI-customer interfaces — is not optional once a chatbot is producing legally consequential statements.

CASE STUDY 05 OF 14

Boom Supersonic

Level 3 (Engine) success pattern: protocols designed alongside capabilities.

Blake Scholl · Move fast and don't break safety-critical ↗

THE PROBLEM

Boom was developing the XB-1 demonstrator under aviation-safety expectations that allow zero tolerance for unexamined design changes. Each iteration on the structural model required stress, fatigue, aerodynamic, and material-tolerance validation before fabrication clearance. Boom needed to accelerate the loop without weakening the verification stack — and to explore hundreds of structural variants in the time conventional methods explore tens.

THE APPROACH

Boom built mkBoom, an internal tool that wrapped AI-assisted structural analysis around the existing verification stack. The AI did not replace any verification step; it accelerated the iteration loop feeding into it. The governance protocols — who owns which analysis, which checks must pass before fabrication, what evidence is preserved for certification — were designed alongside the capability, not added after a near-miss.

THE IMPACT

Boom's XB-1 reached supersonic flight in January 2025 with zero safety incidents across the program. The mkBoom workflow also surfaced Boomless Cruise — a flight regime in which the supersonic boom does not reach the ground under specific atmospheric conditions, recovered through AI-accelerated exploration of the structural envelope. When the blast radius is real hardware, protocols-with-capability is the only viable order: designing them after a failure costs more than the failure itself.

CASE STUDY 06 OF 14

Visicalc & Excel

Level 1 (Shadow) substrate: a vernacular that emerged through pedagogy, not mandate.

Backchannel · A spreadsheet way of knowledge ↗

THE PROBLEM

By the late 1970s, business analysts performed structured calculations on paper, on calculators, or on bespoke mainframe programs maintained by IT departments with weeks-long lead times. The analyst who could think the model was rarely the one able to build it. Visicalc launched in 1979 as the first recalculating spreadsheet — written by Dan Bricklin and Bob Frankston for the Apple II — introducing a grid of cells with formulas referring to other cells.

THE APPROACH

Through three commercial transitions — Visicalc, Lotus 1-2-3, Excel — the file format evolved but the conceptual model did not. Rows, columns, cells, formulas, references: the vocabulary held. No central authority defined or governed it. Practitioners taught each other peer-to-peer; business schools wrote textbooks; managers hired analysts who knew the vernacular without specifying which product. The tool taught the method, and the method outlived three products.

THE IMPACT

By 2026 over a billion people operate through spreadsheet logic that nobody designed and no central authority approved. The L1 Shadow analogue is exact: an immensely productive substrate built by users for users, never sanctioned, never standardized at the top, but absorbed so completely into how work is done that it became invisible. The lesson for AI: substrates of this kind emerge through artifacts that teach the method, not through mandates that prescribe it.

CASE STUDY 07 OF 14

Gates's Internet Tidal Wave memo

Level 2 (Sanctioned) sanction: ratifying cultivated practice already underway.

The Internet Tidal Wave (Gates, 1995) ↗

THE PROBLEM

By early 1995, Microsoft was the dominant force in personal computing, organized around Windows, Office, and a vertically integrated server stack. The Internet — open protocols, browsers, hypertext — was not part of its product roadmap. Inside Microsoft, however, developer interest had been building for over a year: small teams were experimenting with web servers, browser prototypes, and TCP/IP-native tools, mostly outside official product lines. Without executive sanction, the projects could not draw the resources to ship.

THE APPROACH

On May 26, 1995, Bill Gates issued an internal memo titled "The Internet Tidal Wave" to senior staff, declaring "I assign the Internet the highest level of importance." The memo did not invent Microsoft's Internet capability — it ratified work that had been happening unofficially for fifteen months. Attached were specific product directives: a browser shipped with Windows, mail integrated with TCP/IP, a re-architected server stack. The force came from naming what was already happening.

THE IMPACT

Within two years Internet Explorer 1.0 shipped with Windows 95, Outlook and Exchange were repositioned around Internet protocols, and Microsoft was a credible Internet platform company. The case shows what effective Level 2 Sanctioned action looks like: it codifies what cultivation has already proved, channeling existing energy into formal direction. Sanction works when it ratifies; it fails when it mandates without underlying practice.

CASE STUDY 08 OF 14

Git & GitHub pull requests

Level 3 (Engine) protocol: a designed workflow that outgrew its host.

Git kernel docs · BitKeeper usage history ↗

THE PROBLEM

In April 2005, BitKeeper — the proprietary version-control system the Linux kernel team had been using under a free-of-charge license — revoked the license after a reverse-engineering dispute. The kernel team needed a replacement that matched BitKeeper's distributed-workflow properties; nothing on the market did. Beyond Linux, the software industry ran on a patchwork of centralized version-control systems whose mental models did not match how distributed teams actually worked.

THE APPROACH

Linus Torvalds wrote Git in two weeks, optimizing for performance, distributed branching, and verifiable history. Git itself was a powerful but austere substrate — capable of any workflow but prescribing none. Three years later, GitHub introduced the pull request: a branch wrapped in a discussion thread with inline code review, status checks, and eventually CI/CD integration. The PR was not part of Git's specification; it was a designed protocol layered on top — a workflow people could see, comment on, approve, and audit.

THE IMPACT

By 2026 the PR workflow is the assumed professional standard for roughly 100 million developers across essentially every development environment, including ones that do not run on GitHub. The case shows the L3 Engine pattern at scale: a designed coordination protocol on top of a generic substrate, optimized to the point of becoming invisible. It also shows the cultivated-governed handoff working as intended — Git the open substrate, PRs the designed protocol; both necessary, neither sufficient.

CASE STUDY 09 OF 14

Walmart's EDI mandate

Level 4 (Infrastructure) cascade: one big buyer's mandate becoming an industry baseline.

Crstl · Walmart EDI requirements guide ↗

THE PROBLEM

By the mid-1980s, retailers managed thousands of supplier relationships through paper purchase orders, faxed invoices, and phone-confirmed shipment notifications. Error rates at the supplier-retailer interface — wrong quantities, late shipments, duplicate invoices — were a significant fraction of margin. Each retailer had its own document formats; each supplier maintained custom paperwork per customer. Walmart, growing rapidly, found that coordination overhead capped its expansion rate.

THE APPROACH

In 1988, Walmart mandated EDI compliance using the X12 transaction set: adopt it, or lose Walmart's business. The mandate was paired with Walmart's own Bentonville distribution infrastructure, redesigned around the Advance Shipment Notification — a specific X12 message type that allowed receiving docks to plan for incoming shipments before they arrived. Walmart did not invent EDI; it required it, and made compliance valuable to suppliers (faster receiving, fewer disputes) rather than merely costly to refuse.

THE IMPACT

The protocol cascaded across retail within a decade. Suppliers required to use EDI for Walmart began using it elsewhere; competitor retailers added EDI requirements; ERP vendors built EDI tooling into their software. By 2026, roughly 95% of US B2B supply-chain document exchange runs on EDI variants. The case shows the L4 Infrastructure pattern: protocols at this level spread when one dominant actor enforces them against its trading partners, and the rest of the ecosystem follows because the cost of two standards is higher than the cost of one.

CASE STUDY 10 OF 14

McLean's container protocol

Level 5 (Planetary) coordination: an artifact-defined protocol no one owns.

THE PROBLEM

Through the 1950s, freight moved between ship, rail, and truck through "break-bulk" handling — workers physically transferring barrels, crates, bales, and boxes at every transition point. Loading a ship took days; unloading at the destination port took days more. Theft was endemic, damage frequent, and the dwell time of a single shipment measured in weeks. The bottleneck was not transportation but the handoff between modes.

THE APPROACH

On April 26, 1956, trucker Malcolm McLean sailed the SS Ideal X from Newark to Houston with 33 standardized metal boxes — containers — lifted directly from trucks onto the ship and back. The container itself was the protocol: fixed-dimension, lockable, stackable, reusable, designed so cargo never had to be touched between origin and destination. The standard hardened in stages: McLean's company published dimensions; competitors adopted them; ISO 668 codified 20- and 40-foot TEU dimensions in 1968.

THE IMPACT

By 2026 approximately 90% of non-bulk world trade moves via containers. The protocol made coordination cheap and uniform across every jurisdiction and mode. It also concentrated the failures: the system depends on five geographic chokepoints — Suez, Panama, Malacca, Hormuz, Bab el-Mandeb — whose vulnerability was made visible by events like the 2021 Suez Canal blockade. The L5 Planetary lesson: when a protocol becomes infrastructure for the world, coordination is cheap and legibility becomes the new governance challenge.

CASE STUDY 11 OF 14

Rural users and the Model T

Level 1 (Shadow): when users invent uses the manufacturer never imagined.

Kline & Pinch · Users as Agents of Technological Change (1996) ↗

THE PROBLEM

When Ford launched the Model T in 1908, it was designed as urban transportation. Rural America had a different problem: small farms needed power for threshing, sawing, water-pumping, washing, and churning. Buying a dedicated stationary engine for each task was prohibitively expensive. The technology was widely available; the use case that mattered to half of Ford's customers wasn't visible to the manufacturer.

THE APPROACH

Rural users improvised. Owners jacked Model Ts onto blocks and ran belts off the rear wheels to drive threshing machines, washing machines, water pumps, fodder cutters, and circular saws. Third-party accessory makers shipped kits (the "Stationary Hoister"); farm magazines documented dozens of unauthorized uses by the 1920s. Ford initially resisted, then gradually acknowledged the use cases in marketing materials and parts catalogs, redesigning later models to support them.

THE IMPACT

The kit-phase pattern Kline & Pinch document is structural: user adaptation precedes manufacturer recognition by years. Value is generated by users in ways the supplier doesn't anticipate; the supplier's eventual job is to harvest and standardize what users invented. The same pattern is operating today around AI — practitioners are using AI for purposes their employers haven't sanctioned. The organizational job is to surface those uses (Kit Harvesting), not suppress them.

CASE STUDY 12 OF 14

American steam adoption, 1820–1900

Level 2 → Level 3: firms make the architectural commitment to a new paradigm for what it enables, not for what it makes cheaper.

Atack, Bateman, & Weiss · The Regional Diffusion and Adoption of the Steam Engine in American Manufacturing (1980) ↗

THE PROBLEM

American factories in the 19th century had a recurring choice between two prime movers. Water power had been the standard since the early industrial period — abundant, well understood, and, in most parts of the country, measurably cheaper per horsepower than steam. Cost simulations by Atack, Bateman, and Weiss show that water was the lower-cost option in New England, the Middle Atlantic, and the South for most of the century. Yet steam adoption climbed steadily anyway. By 1900, steam engines outnumbered water sources four to one. The cost ledger and the adoption pattern were pointing in opposite directions, and the gap held for decades.

THE APPROACH

Atack and his coauthors traced what firms actually chose, region by region, between the 1820s and the 1890s. The pull toward steam wasn't operating cost — it was locational freedom. A water-powered factory had to sit by the water; a steam-powered one could sit anywhere, including next to railroads and urban labor pools. Pittsburgh and Cincinnati adopted steam earlier than Boston because the eastern rivers were already worked. Augusta, Georgia, had ample water power but was over 100 rail miles from major urban markets, and never became a manufacturing center. Having the input wasn't enough.

THE IMPACT

American manufacturers paid a premium for steam for most of a century. They were paying for capabilities the cost-per-horsepower table didn't capture: where a factory could sit, how large it could grow, who it could hire, how it could ship. There's a recognizable echo here for organizations evaluating AI adoption now. When teams compare AI's cost to the cost of the work it might replace, the math often looks marginal or worse, and the natural conclusion is to wait. The historical pattern suggests a complementary question worth asking alongside the cost one: what becomes newly possible? Work that was previously out of economic reach, problems that couldn't be attempted at all, customers or markets that weren't addressable at the firm's prior scale. Adoption often makes sense when the answer to that question is concrete, even in cases where the cost comparison on its own doesn't make the case.

CASE STUDY 13 OF 14

Factory electrification, 1880–1930

Level 2 → Level 3: the 30-year lag between substitution and architectural redesign.

Warren D. Devine Jr. · From Shafts to Wires (1983) ↗

THE PROBLEM

Through the 19th century, factories were powered by central steam engines driving overhead line shafts and belt drives to each machine. The architecture imposed constraints: machine placement followed shaft geometry, multi-story buildings minimized shaft runs, whole-plant on/off operation, 25–40% friction losses, blocked windows from overhead shafts. When electric motors became commercially available in the 1880s, the obvious move was to swap the prime mover.

THE APPROACH

Devine documents three phases. (1) Group drive (1890s–1910s): the steam engine was replaced with a large electric motor, but the line shaft, belts, and machine placement remained unchanged. By 1919, 53% of US installed manufacturing horsepower was electric. (2) Unit drive (1910s onward): a dedicated motor at each machine, sized to its load, with wiring replacing the iron shaft. Factory layout finally followed workflow, not shaft geometry. (3) Architectural redesign: single-story buildings (Ford's River Rouge, 1928), sawtooth roofs for daylight, integrated conveyors, 40–50% reductions in energy per unit output.

THE IMPACT

US manufacturing TFP grew under 1% per year through the group-drive era (1899–1919) and ~5% per year in the unit-drive era (1919–1929). The productivity boom of the 1920s was the payoff of architectural redesign, not of electrification per se. Substituting the power source without redesigning the layout captured almost none of the available gains. Most current AI deployments are at the equivalent of group drive — the technology is installed, the workflow is unchanged, and the productivity hasn't materialized. The L2 → L3 transition is what closes the gap.

CASE STUDY 14 OF 14

Harpers Ferry and the American System

Level 3 → Level 4: the institutional substrate behind "interchangeable parts."

Roger D. Simon · The Machine in Context (review of Merritt Roe Smith) ↗

THE PROBLEM

Eli Whitney's 1798 federal contract for muskets promised interchangeable parts — pull any lock, any barrel, any stock from a pile and they fit. The marketing claim was repeated through the early 19th century. The actual practice was different: parts were fitted individually by skilled file-and-fitter armorers, and each musket was effectively bespoke. The technology label existed; the operating practice didn't.

THE APPROACH

Merritt Roe Smith's history of Harpers Ferry Armory (1798–1860s) documents a long institutional development. Federal armory contracts created a sustained customer that insisted on uniformity. Machine tools were developed in-house (milling machines, pattern lathes, gauges). Inspection protocols co-evolved with the machines. Workers were retrained — gauge-using inspectors replaced file-and-fitters. The social order of the armory changed alongside the technical order. By the 1850s, the American System of Manufactures could be demonstrated at the London Crystal Palace Exhibition, three decades after the rhetoric began.

THE IMPACT

Industry-scale uniformity required the full institutional substrate to mature, not just the technical capability. The same is operative for L4 Infrastructure today: stating a standard isn't the same as having the inspection regimes, training systems, and shared assumptions that make the standard actually hold. Organizations claiming AI standardization without that institutional substrate are repeating Whitney's marketing claim, not Harpers Ferry's actual practice.

Every node and every edge of how work moves is now AI-mediated

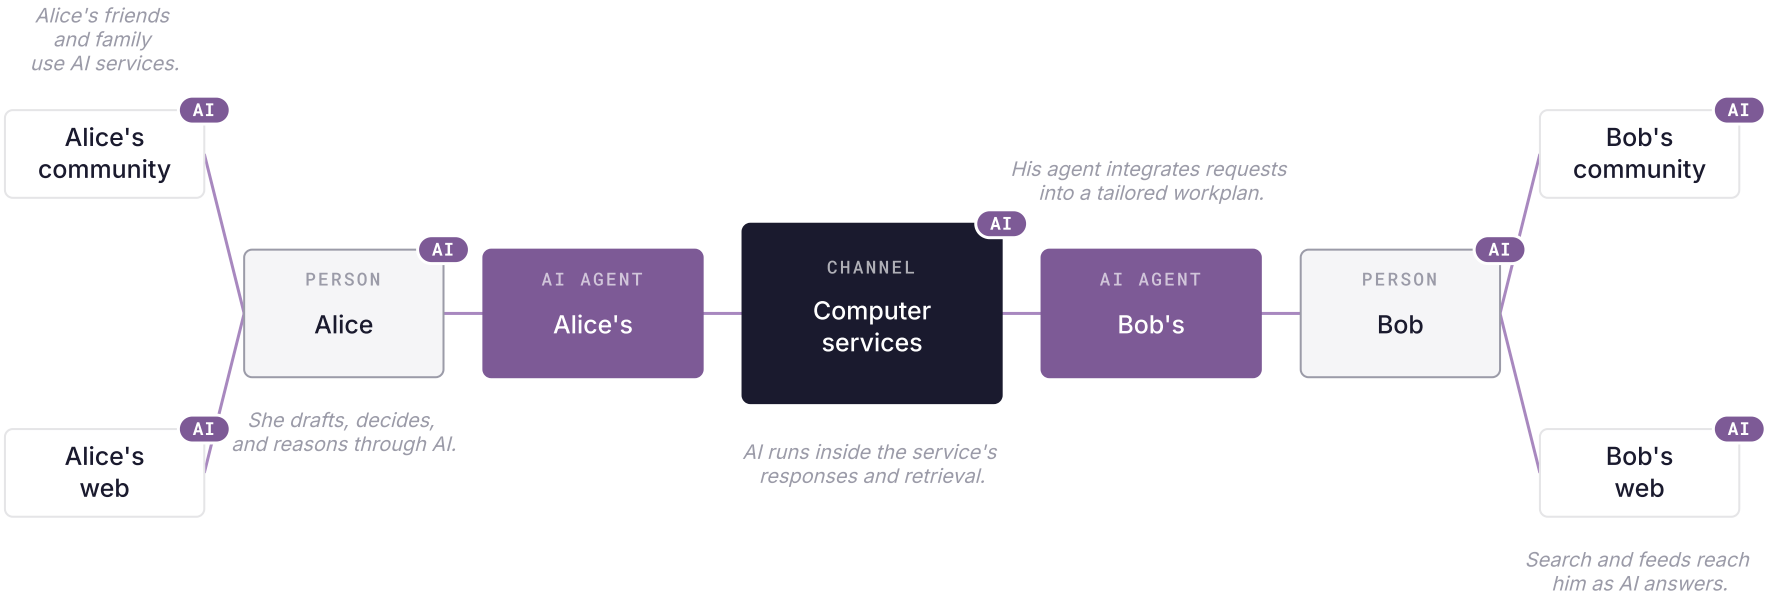

Finding the right artifact for a given technology

The artifact that matches the framing. AI sits in the Paradigm × Planetary quadrant — the artifact that fits it is an Adoption Guide.

	✗ LOCAL	✓ PLANETARY
✓ PARADIGM	**Practice Playbook** e.g., MRI AI: Vibe-Coding	**Adoption Guide** e.g., electricity, the internet **AI: Commoditized Generative Software**
✗ TOOL	**Implementation Roadmap** e.g., Figma AI: Helpdesk Chatbot	**Deployment Plan** e.g., Microsoft Office AI: Claude Code

Where is your organization?

A 2-minute diagnostic that places you on the maturity ladder and identifies your dominant archetype.

QUESTION 1 OF 25

What scope are you answering for?

- ○ A small organization (under 200 people)
- ○ A business unit or division of a larger enterprise
- ○ A specific function (engineering, product, operations, sales, support, etc.)
- ○ A specific team or workstream