

The System of Work

Christian Blank · May 2026

We are at the beginning of the largest leap in collective effectiveness in human history. AI does not just make individuals more productive — it changes who can participate, how fast, and at what scale. The catalyst that turns this potential into compounding reality is not more capability. It is the coordination layer that connects every actor to shared truth, shared intent, and shared governance. This document is about that layer.

The Thesis

Every organization that ships anything runs on a coordination substrate. Sprint plans, org charts, culture handbooks, release processes, on-call schedules, communication etiquette, document templates, board notes, operator manuals, road maps, retrospective formats, approval chains, onboarding checklists — all of it is the system of work. Scattered across tools, habits, documents, and social memory. Some parts conscious, most unconscious. Some parts strict, most implicit. High entropy. Low coherence. Functional enough because humans perform constant quiet repair to fill the gaps.

This substrate is not new. It predates software. It predates computing. Wherever multiple actors coordinate toward shared purpose — a company, a team, a family, a city, a volunteer build — the substrate is there, carrying the load. The principles that govern it are old and stable: shared truth, governed authority, traceable provenance, aligned intent, evolving feedback. These principles do not depend on who the actors are. They depend on the fact that participation is plural.

The System of Work Is Not an AI Thesis

The System of Work matters with or without AI.

Any organization that works well already has some effective version of this layer. Manufacturing had Lean. Software had Agile. Operations had SRE. Infrastructure had DevOps. Product has discovery, planning, review, launch, and iteration loops. Research has papers, experiments, evals, peer review, and replication. Customer-facing teams have

escalation paths, account notes, success plans, qualification frameworks, and feedback loops.

Each of these practices solves part of the same problem. They make work more observable, repeatable, governable, and improvable inside a local domain. But the domains do not stay separate. Research becomes product. Product becomes platform. Platform becomes infrastructure. Infrastructure becomes operations. Operations produces customer signals. Customer signals reshape product. Product constraints reshape research. Research constraints reshape infrastructure.

Local frameworks help, but they are not the unified substrate. Agile does not solve architectural memory. Lean does not solve AI-agent governance. SRE does not solve product intent transfer. A design doc does not solve customer feedback compounding. An ADR does not solve active work coordination. A project management tool does not solve provenance. A knowledge base does not solve authority. A chat tool does not solve organizational memory.

Each tool or framework handles a slice. The System of Work is the shared layer underneath the slices.

It asks: Who is acting? Against what truth? With what authority? Toward what intent? Using what memory? In coordination with whom? Based on what provenance? With what capability? Under what constraints? In service of what mission?

That is not an AI problem. That is a work problem. AI makes the answer more urgent because AI increases the speed, scale, and number of actors operating inside the system. But the need is older than AI. The substrate has always been there.

Three Claims

The substrate is a distinct layer — and it has always evolved, not been designed. It has its own primitives, its own dynamics, its own evolution curve. Most organizations leave it implicit — scattered, fractured, carried by tenure and social repair. The question is not whether to design the substrate. It is whether to design the conditions under which the substrate evolves. Building that meta-system explicitly compounds. Leaving it implicit decays.

AI is the catalyst that made this urgent, not the reason the substrate exists. AI agents do not perform the social repair that kept the implicit version functional. They do not infer authority from social context. They do not know what was decided last quarter unless it is encoded somewhere they can read. When AI enters the coordination structure, every gap in the substrate becomes audible. The substrate was always load-bearing. AI is what made it

impossible to ignore.

Every improvement in model capability increases the demand for this layer. Faster models mean more concurrent actors. More concurrent actors mean more shared-state mutation per unit time. Smarter models mean higher-stakes autonomous action, which requires explicit governance. The coordination demand grows with capability. The leverage of building the meta-system that lets the substrate evolve grows with it.

The Collaboration Problem

Most valuable work is not isolated. It crosses boundaries. Two teams need to touch the same architecture. A product surface depends on a platform primitive. A research capability needs to become production infrastructure. An internal tool starts as a local workaround and becomes shared infrastructure. A customer request reveals a product gap that touches three domains.

Today, that creates alignment overhead. Someone has to understand the current architecture. Someone has to know the history. Someone has to remember why previous decisions were made. Someone has to know which ADRs are still valid. Someone has to know which specs drifted from the code. Someone has to know which future changes are already planned. Someone has to know which constraints are real and which ones are obsolete.

That person is usually the architecture owner, senior engineer, EM, PM, designer, customer lead, or operator who carries the system in their head. They become the human reconciliation layer. They watch active work. They watch proposed work. They detect conflicts. They interpret intent. They explain history. They review proposals. They run alignment meetings. They keep teams from duplicating effort. They prevent local decisions from degrading the global system.

That role is valuable. It also becomes a bottleneck.

The organization has two bad choices.

Move fast without alignment, and the system drifts. Code duplication rises. Teams build their own versions because it is faster than coordinating. Architecture becomes incoherent. Product patterns fragment. Customer-specific workarounds accumulate. The shared platform becomes harder to change.

Or require heavy alignment, and velocity drops. People wait for meetings, proposals, reviews, sign-offs, context downloads, and political agreement. Lead time increases. Contribution becomes expensive. Eventually, people stop trying to improve shared systems

because the coordination cost is too high.

Both outcomes are system-of-work failures. One creates entropy. The other creates fatigue.

The problem is not that alignment exists. Alignment is necessary. The problem is that alignment is reconstructed manually from scattered memory every time meaningful work crosses a boundary. That is the part that does not scale.

What the System of Work Changes

A System of Work does not remove alignment. It changes the cost structure of alignment.

If context, history, reasoning, decisions, constraints, ownership, active work, future work, and provenance are explicitly preserved, then a human or AI actor can reason from the substrate instead of starting from social discovery. The first proposal gets better. The first spec gets closer. The first implementation fits the architecture more often. The first customer solution compounds more often.

Humans stop being the blue team for every proposal. They become the red team. Their job shifts from generating and re-explaining the whole context to challenging, correcting, validating, and improving a near-first-pass artifact.

That is the leverage. The organization does not need every contributor to rediscover the architecture before contributing. It needs the architecture, its history, its product reasoning, its ownership model, its customer context, and its evolution rules to be represented well enough that contributors can enter the system with minimum coordination overhead.

The substrate should help any actor — human or AI — answer: What is the current state? Why is it that way? Who owns it? What decisions shaped it? Which constraints still matter? What future work is already planned? Where would this proposed change conflict? What pattern should this follow? What should this reuse? What would make this safe to approve? What is the cleanest next move?

Today, those answers are scattered across people, docs, code, specs, ADRs, tickets, chat, meetings, and memory. A System of Work makes them part of the operating layer.

Capability Transfer

A System of Work does more than preserve context. It transfers capability.

In most organizations, capability is trapped inside roles, people, and teams. Product judgment lives with product managers. UX taste lives with designers. Architectural

reasoning lives with senior engineers. Customer context lives with sales and customer success. Operational judgment lives with support and infrastructure teams. Historical knowledge lives with whoever happened to be there.

This creates unnecessary dependency at every boundary. An engineer who sees the product implication of a technical choice still has to route through product. A product manager who understands the customer need still cannot safely implement or prototype the feature. A customer success manager who sees the same pattern across five customers still has to negotiate with multiple product owners before that pattern can become a compounding product improvement. A designer who understands the user problem may not know which system constraints are real. An operator who sees a workflow gap may not know how to convert it into a shared tool.

When the substrate is explicit, that changes. The system can expose the current product intent, UX patterns, architectural boundaries, customer reasoning, ownership model, accepted tradeoffs, active constraints, and provenance behind prior decisions. That lets more people contribute correctly.

This does not erase roles. It makes roles more leverageable. People still bring different judgment, taste, authority, and accountability. But they no longer need every adjacent domain expert to manually reconstruct the context before meaningful contribution can begin. The system carries more of the organization's accumulated judgment. The human applies more of their own judgment.

That is the difference between role-bound execution and substrate-enabled contribution. A System of Work turns expertise from a private dependency into shared infrastructure. The point is not that everyone can do everything. The point is that more people can do the right thing without first extracting the organization from someone else's head. Roles remain. Bottlenecks shrink.

The AI-Native Version

AI makes this more urgent because AI can generate work faster than organizations can align around it.

A coding agent can produce an implementation. A planning agent can produce a spec. A product agent can produce a workflow. A platform agent can propose an abstraction. A research agent can suggest a deployment path. A support agent can summarize customer pain. A design agent can produce flows. A data agent can surface patterns.

But without explicit substrate, the output is only locally intelligent. It may not know the architectural history. It may not know the ownership boundary. It may not know the future

migration. It may not know why the obvious solution was rejected. It may not know which team is already changing the same surface. It may not know the UX pattern the product is trying to standardize. It may not know which customer workaround should become product and which should remain custom. It may not know whether the request conflicts with governance.

So humans become the integration layer again — reviewing not only quality, but context. Not only implementation, but fit. Not only correctness, but architectural coherence. Not only customer value, but product direction. Not only speed, but governance. That does not scale.

The better version is a domain-aware System of Work. The system preserves the context, provenance, decisions, ownership, constraints, active work, future work, and observed outcomes of a domain. Over time, it becomes capable of producing better first-pass proposals because it is reasoning against the actual substrate, not a disconnected prompt.

Eventually, a domain agent can become the first gatekeeper for a slice of the architecture, product, workflow, or customer domain — not the final authority, but the first coherence check. It can say: This proposal conflicts with an active migration. This repeats a rejected pattern. This changes a contract owned by another domain. This creates a duplicate tool instead of extending the shared primitive. This needs a governance review. This matches the current product direction. This should reuse an existing UX pattern. This is compatible with the current architecture. This is safe as a local extension but not as a platform abstraction yet. This customer request reveals a broader product primitive. This spec is missing the reasoning needed for future contributors.

That is not just AI assistance. That is substrate-aware collaboration.

The Arc

The system of work has always evolved. Humans did not design a system to coordinate. When we became more conscious, more ambitious, more capable, more connected, we started building frameworks — not because we invented coordination, but because the gap between our potential and our current effectiveness became visible enough to close on purpose.

In early human organization, coordination was implicit and emergent. Small groups organized around survival. The substrate evolved at the pace of the environment — slowly, through proximity, through high-frequency face-to-face interaction that naturally smoothed friction and propagated knowledge. Implicit was fine because the cadence was slow and the actors were few.

As groups grew, we specialized. Some produced food. Some figured out health. Some built shelter. Some created governance. The larger the group, the more explicit the coordination had to become — not because the principles changed, but because implicit repair does not scale. A village can coordinate through social memory. A city cannot. A nation cannot. Each scale transition forced more of the substrate into explicit form — laws, institutions, roles, processes, records.

Industrialization compressed the timeline. The gap between what was possible with the new machines and what organizations were actually achieving became visible — and expensive. So we formalized it. Organizational psychology, sociology, management theory. Frameworks: Kanban, Six Sigma, Lean, Agile. Each one was an attempt to close the gap between potential and practice by making part of the substrate explicit and repeatable. Each one worked for its era. Each one eventually dated as the environment changed.

Computers compressed the timeline further. A person with a computer was ten times more productive than one without. But the human remained the interface — the bottleneck through which all work flowed. The asymmetry between human speed and machine speed was manageable because the machine was a tool, not a participant. The machine did not make decisions. It did not hold intent. It did not coordinate with other machines toward a shared mission. The substrate could remain mostly implicit because the only actors who needed to coordinate were still human.

AI changes this in kind, not in degree. For the first time, we built a technology modeled after the human brain — a generalist that becomes a specialist on demand, that scales with compute rather than with twenty-five years of human upbringing. It is non-deterministic in the way humans are non-deterministic. No simulation captures all its states. No test suite covers all its behaviors. And unlike every prior machine, it does not stay in the tool lane. It participates. It holds intent across time. It takes actions with real consequences. It coordinates with other actors — human and AI — inside shared structures.

The closest precedent is the self-driving car. A non-deterministic actor granted first-class citizenship in a high-stakes environment, operating alongside humans who did not choose to interact with it. The environment has effectively unlimited parameters. Monitoring and behavior prediction are close to impossible. And yet we granted this actor the same authority to make consequential decisions as any human participant. The only way to not interact with it is to retract from public roads.

The trajectory for AI in work is the same. The difference between a coding agent and a self-driving car is that one has a physical body and the other a digital one. The capabilities differ. The level of autonomy and authority being granted will converge. Both participate alongside other actors toward shared missions. Both contribute autonomously toward milestones based on their role. The only way to not interact with AI actors in the future of work is to retract from professional life.

And now the asymmetry accelerates beyond anything we have experienced. Computers compressed the cycle by an order of magnitude. AI compresses it by two orders. But humans cannot evolve at that speed. We cannot absorb, adapt, and develop new coordination patterns a hundred times faster than we did before. The gap between potential and practice widens faster than any individual or organization can close it through natural emergence.

The response so far has been tools — better monitoring, better developer tools, better ways to check AI output, correct it, steer it. These tools are useful. They are also insufficient. They treat the symptom — the human struggling to keep up with the machine — without addressing the structure. They optimize the human's ability to supervise AI within the existing implicit substrate, rather than building the substrate that lets humans and AI coordinate as equals.

A hard-coded self-driving car does not work because the environment has unlimited parameters and evolves continuously. A hard-coded system of work does not work for the same reason. You cannot design the substrate and deploy it. You have to design the system that enables the substrate to emerge, evolve, calibrate, and adapt — at the cadence the environment now demands.

This is the System of Work. Not a designed framework. A designed capacity for emergence. The meta-system that lets the coordination substrate evolve through use, at the speed of its most capable actors, governed by the actors who are accountable for the outcomes.

What Became Visible

The gaps that AI exposes are not new. They are gaps that humans papered over through social repair, tenure, and institutional memory that lived in people's heads.

What felt like culture turns out to have been infrastructure — implicit rules about who communicates what to whom, carried by habit rather than by system.

What felt like seniority turns out to have been a context cache — years of accumulated knowledge about how things work here, never written down, lost every time someone left.

What felt like trust turns out to have been a lookup table of authority — who gets to approve what, who owns which domain, who has the final word — carried in social memory rather than in an inspectable system.

Every organization experiences this. A new hire takes months to become productive not because the work is hard but because the substrate is implicit and can only be absorbed through exposure. A reorganization resets institutional knowledge not because the people

changed but because the social memory that carried the substrate scattered. A process that works perfectly for five people collapses at fifty not because the process was wrong but because the implicit repairs that made it work do not scale.

AI compressed this problem by operating at a cadence and concurrency that implicit repair cannot match. But the problem was always there. Any organization that has ever lost institutional knowledge when someone left, or struggled to onboard new people, or watched a process break when the team grew — that organization was experiencing the cost of an implicit substrate.

The opportunity is not to build something new for AI. It is to build the meta-system that lets the substrate evolve explicitly — at the cadence the environment now demands, for any actor at any scale.

The Primitives

The substrate decomposes into seven primitives. They are not invented. They are discovered repeatedly in any sufficiently hard coordination problem. Every organization exhibits all seven, whether or not it has named them.

Identity. The inherent nature of the actor — not a label the system assigns, but what the actor actually is. Constants that cannot change and capabilities shaped through experience. An AI model has an architecture, capability boundaries, and training biases it cannot override. A human has a cognitive style, knowledge, and judgment patterns — some inherent, some developed. A penguin in the desert is still a penguin. Trust, credibility, and standing are identity as perceived by others — earned through outcomes, lost through failure, always evolving. In an implicit substrate, identity is reputation and tenure. In an explicit one, it is a structural, observable record of who the actor actually is.

Context. The shared truth all actors operate against right now. Without explicit context, every actor assembles their own private version of reality, and the system drifts apart at the speed of its parallelism. In an implicit substrate, context lives in the heads of long-tenured people. In an explicit one, context is a queryable, shared, maintained artifact.

Memory. What persists across time. The record of what has been observed, decided, tried, learned, and rejected — available across sessions, across actors, across team changes. Memory is not logging. It is the longitudinal substrate that lets an organization know its own history. Most organizations lose memory every time someone leaves. An explicit substrate retains it structurally.

Coordination. How concurrent action stays consistent. Multiple actors working simultaneously on overlapping concerns need explicit semantics for how their work

composes, how conflicts are detected, and how they are reconciled. This is a distributed-systems problem whether the actors are humans in different time zones, AI agents in parallel threads, or automated pipelines running concurrently.

Intent. How purpose propagates and is verified. Every effort carries an intent — what this is for, whose mission it serves, what good looks like. Intent has to travel with the work, be readable by every actor, and be verifiable against what actually gets done. In an implicit substrate, intent lives in the head of whoever started the project. In an explicit one, intent is a structural artifact that governs downstream work. Intent ungoverned drifts. Intent encoded compounds.

Governance. How actors are authorized to operate — the full spectrum from explicit to implicit. Explicit governance is processes, approval chains, release gates. Implicit governance is social capital, credibility, standing — the authority that accrues through demonstrated judgment and erodes through poor outcomes. Both forms are always present. Both evolve continuously. The signals are continuous; the organizational responses are discrete — promotions, scope changes, demotions. The gap between continuous signal and discrete response is where most organizational friction lives. In an implicit substrate, these two forms are entangled and political. In an explicit one, the relationship between evidence and authority change is inspectable.

Provenance. What happened, why, and on what basis. Every artifact, decision, and recommendation carries the trace of its origin — the observations that produced it, the reasoning that shaped it, the outcomes that confirmed or revised it. Provenance is how the system evaluates, evolves, and trusts its own knowledge. Without it, decisions persist by inertia rather than by validity. Without provenance, organizations accumulate zombie knowledge — old decisions that outlive their basis, documents that remain trusted after they become stale, processes that continue because nobody remembers why they were created.

The opportunity is to make these seven primitives explicit, build them as first-class properties of the system rather than emergent properties of social memory, and let them compound across every effort the system runs.

Why It Compounds

An implicit substrate decays. People leave and take context with them. Teams change and social memory scatters. Processes that worked at one scale break at the next. Every reorganization partially resets institutional knowledge. The organization's capability is always one departure away from regression.

An explicit substrate compounds. Every decision is recorded with provenance. Every

workflow encodes what was learned from the last execution. Every governance evolution is structural and permanent. People can leave. Teams can change. The substrate retains everything it has learned.

The compounding accelerates because the substrate improves itself through use. Every execution produces signals about the substrate's own fitness — where a workflow was too tight, where governance was too loose, where context was missing, where intent drifted from outcome. These signals feed back into the substrate. The next execution is better governed, better scoped, and better informed. The system does not just improve. The rate of improvement itself accelerates.

This compounding applies to any organization, regardless of whether it uses AI. A ten-person team with an explicit substrate onboards new people faster, retains knowledge across departures, and evolves its processes from evidence rather than from opinion. A hundred-person company with an explicit substrate scales coordination without the communication overhead that normally grows quadratically with headcount.

AI amplifies the compounding because AI increases the cadence. More concurrent actors produce more executions per unit time. More executions produce more signals. More signals produce faster evolution. The substrate improves at the speed of its use, and AI uses it faster than humans do. This is why model capability and substrate value move in the same direction — smarter, faster models extract more value from better context, and the substrate is what produces the context.

There is a historical pattern. Compute needed an operating system. The internet needed protocols. Cloud needed orchestration. Each capability layer required a coordination layer to turn raw power into a composable environment. AI-native work needs a system layer. The leverage accrues to the layer that turns participation into coordination. That is the substrate.

What This Enables

A mature System of Work changes what an organization can do.

It reduces alignment overhead without removing alignment. It increases autonomy without abandoning governance. It improves velocity without accepting entropy. It transfers capability without erasing roles. It lets more people contribute without making the system incoherent. It gives AI agents enough substrate to participate without forcing humans to become full-time integration layers.

The practical effect is concrete. An engineer can make better product decisions because product reasoning is visible. A product manager can prototype or implement safely because architectural constraints and contribution paths are visible. A customer success manager

can explain what a feature is optimizing for and propose a compounding product solution instead of negotiating a one-off workaround. A designer can work with real system constraints instead of guessing what engineering will reject. An operator can turn repeated internal pain into a shared tool that follows the organization's patterns. A researcher can see how a capability moves toward product, infrastructure, evals, safety, deployment, and operations. An AI agent can generate a spec that starts from the organization's actual context instead of from a prompt detached from history. A domain owner can review higher-quality first-pass work instead of restating the same background in every alignment meeting.

The organization becomes easier to contribute to — not because the work is simpler, but because the substrate carries more of the complexity.

The Final Claim

The future of work is not just faster individuals. It is more coherent participation. More humans. More AI agents. More tools. More automation. More domains. More parallel work. More shared-state mutation. More need for governance. More need for context. More need for memory. More need for provenance.

The organizations that win will not be the ones that simply add AI to existing implicit systems. They will be the ones that make the operating substrate of work explicit. They will know what the organization knows. They will preserve why decisions were made. They will let people contribute across boundaries without starting from social discovery. They will let AI agents participate inside governed context instead of disconnected prompts. They will make expertise transferable. They will make alignment cheaper. They will make governance inspectable. They will make execution compound.

The substrate has always been there. The opportunity now is to build it deliberately.

— Christian