
THE FOUNDING ARGUMENT · 2026

AI Operations Management

The Discipline of the Flow

Daniel S. Wipert

On the economics of business AI adoption

The category error

Enterprise software has long been managed on a defining economic assumption: that its marginal cost of use is effectively zero. Once licensed and installed, the hundredth use of a spreadsheet application costs the organization no more than the first. Nearly every management structure surrounding enterprise technology, from procurement and budgeting to cost accounting and vendor governance, is built on this assumption, whether explicitly or by habit.

Deployed AI does not have this property. Each use of an AI system is a consumption event. Computational resources are expended, and the expenditure is measurable, meterable, and attributable to the activity that caused it. A single contract review, a single drafted reply email, a single classified invoice: each is such an event, and a firm that has deployed AI broadly runs them by the thousands before lunch. Usage that repeats consumes repeatedly. Usage that scales consumes at scale. In economic terms, AI is not a fixed-cost asset that an organization holds. It is a variable-cost activity that an organization runs on a daily basis.

The pricing on the invoice does not change this. An organization paying a flat rate for AI access may conclude that its consumption is economically irrelevant, and that the subscription has converted a variable cost back into a fixed one. It has not. Flat pricing conceals the consumption; it does not eliminate it. The underlying resource expenditure remains variable, remains real, and is borne in the first instance by the AI provider, whose own economics are anything but flat. An AI provider carrying variable costs against fixed revenue faces a structural mismatch that will not last. It resolves through meters, usage tiers, rate limits, throttling, priority access, or repricing. Every one of these mechanisms is already visible in the AI market. The price of a unit of AI capacity already moves from one day to the next. An organization that plans around today's flat rate is planning around a pricing artifact rather than an economic reality, and the correction will arrive on the AI provider's terms.

This distinction is not an accounting technicality. It changes the category of the thing being managed. An organization that deploys AI across its workforce and workflows is not administering a software product. It is operating a flow. “Flow”

here is meant strictly, and it will carry the rest of this document: a continuous, high-volume stream of consumption events moving through business processes, each event measurable, each carrying cost by default, each producing value only by design.

Business has long experience with such flows. Goods flow through firms, and money flows through firms, and the consistent lesson of the twentieth century is that an unmanaged flow will always degrade. Costs become invisible, accountability becomes unassignable, and control is lost precisely where volume is highest. This is why the flow of goods is governed by a mature discipline, supply chain management, with defined functions, professional standards, and executive ownership.

The flow of AI usage has no such discipline. The absence is not a failure of diligence. It is the absence of a field.

This document defines that field. Its name is AI Operations Management.

02

The historical rhyme

Every organization deploying AI at scale is now running one of the largest operational flows in its business, and almost none of them manage it as one. AI spending sits in an IT budget line. Usage data sits in provider dashboards that no one consolidates. Value claims sit in vendor presentations. Responsibility sits nowhere. The consumption is real, the functions that touch it are disconnected, and no discipline unifies them.

Business has seen this exact configuration before. Firms moved goods for a century before anyone managed the movement: purchasing, warehousing, and freight each optimized its own corner, and no one owned the whole process flow, because no one had yet conceived of it as a whole. Then scale arrived through containerization, global sourcing, and exploding product variety, and the leakage from the unmanaged flow stopped being tolerable. The response was not a technology. It was a discipline. Supply chain management named the fragmented functions as segments of a single flow, placed that flow under coherent

measurement, and raised ownership of it to the executive level. The goods had been flowing all along. The discipline was the innovation that created order out of the spiraling chaos.

The principle beneath the history is strict. Flows precede their disciplines. A flow can run unmanaged while its volume is modest because at modest volume the leakage is absorbed as noise. When the flow scales, the leakage compounds until a discipline is founded to see it as a whole, measure it, and own it. Management disciplines are not born from theory. They are born from scale, which creates information and cost entropy.

AI usage has crossed that threshold. An organization running scattered pilots can treat consumption as noise, for the same reason a small workshop never needed a logistics department. An organization running AI in production cannot. At production scale, consumption events number in the millions, cost exposure accumulates across every workflow the flow touches, and each unmeasured period widens the gap between what is spent and what is known. This is not a warning about the future. It is a structural claim about the present: scaled AI deployment requires cost governance, because without it the organization has surrendered visibility, attribution, and allocation over one of its largest flows and cost centers. What cannot be seen, cannot be controlled, and most organizations simply do not see it.

There is a second pressure in the pattern, and it is the less comfortable one. Disciplines also arise defensively, because the other side of the transaction organizes first. Firms professionalized purchasing in part because suppliers had professionalized selling. The same asymmetry is forming now, and faster. AI providers carry real costs for every unit of service they deliver and operate under genuine capacity constraints, which gives them a durable economic basis for managing access through meters, tiers, rate limits, and priority. The supply side of the AI market is professionalizing its side of the flow. A buyer facing a managed supplier without a management discipline of its own is not a counterparty. It is a price-taker, and it will be priced accordingly.

History does not repeat its particulars. It repeats its structure. Every flow that reached scale summoned a discipline, and every summons was answered. Firms that organized early were able to answer deliberately and took the advantage.

Firms that waited were exploited severely by the open market. AI usage will not be the exception. An organization does not choose whether to answer. It chooses only when, and the price rises while firms balk.

03

The triple flow

Supply chain management earned its authority from a single act of perception: it saw that a firm does not run one flow but three. Goods move through the firm. Information about the goods moves alongside them. Money moves in the opposite direction. The discipline's founding insight was that these flows are inseparable, that the goods cannot be managed without the information, and that the money cannot be controlled without both. Every function the field later developed, from demand planning to inventory control to procurement, is an instrument for keeping the three flows aligned.

AI Operations Management builds from the same perspective. An organization deploying AI at scale runs three flows, whether it manages them or not, and the discipline that manages them resolves into five functions: sourcing, planning and budgeting, metering and attribution, allocation and routing, and value-boundary management.

The first flow is usage, and it is the primary flow. It begins on the supply side, where capacity to perform AI work is sourced from providers who consume real resources to deliver it. It moves into the organization through deployment, where AI use becomes part of operating activity. AI runs through workflows, where each task performed is a discrete consumption event. Workflows that integrate AI come to depend on it, so that the flow becomes load-bearing for the business processes it passes through. And it terminates in work products: drafts, analyses, reports, classifications, and decisions that are the flow's entire purpose. Source, convert, deliver. This is a supply chain in the strict sense, not a metaphor borrowed for effect. The only unusual property of this supply chain is that the goods are invisible, which is precisely why it has gone mostly unmanaged.

The second is the flow of records, and it shadows the first. Every usage event is measurable, and measurement produces records: what ran, where, initiated by

whom, and consuming how much. Records accumulate into visibility, visibility into the ability to compare, attribute, and decide. This flow is to AI operations what the information flow was to logistics, and the twentieth century's verdict on that flow was unambiguous. The barcode and the shipping manifest mattered as much as the truck, because a goods flow without an information flow is not a supply chain. It is blank sailing in the dark. The same verdict applies here, and it applies with force in both directions. Where usage is metered, records exist and management is possible. Where usage is unmetered, the activity does not merely go unmanaged; it goes unseen, and what is invisible to the management boundary is beyond economic control. The record flow is therefore not administrative overhead attached to the usage flow. It is the precondition for managing it at all.

The third is the flow of cost and value, and it is where the AI Operations discipline earns its keep. Cost enters the flow automatically. Every usage event with a positive price or resource cost creates exposure at the moment it runs. Exposure accumulates across repeated events without any decision being made, and scaled deployment multiplies the accumulation across every user and workflow it reaches. No meeting is held, no approval is sought. Cost simply travels with the flow, as freight charges travel with cargo.

Value does not behave this way, and the difference between the two behaviors is the central economic fact of AI Operations. Value does not enter the flow automatically at any point. A work product delivered is not value realized; a productivity story told is not a return demonstrated. Value exists only where a specific benefit is actually captured inside a defined boundary: this workflow, this period, this measured outcome, net of the full cost of producing it. Claimed value is abundant in the AI economy. Realized value must be extracted from the flow deliberately, at a boundary someone has drawn, by someone accountable for drawing it.

This is the asymmetry on which the entire discipline stands. *Cost accrues by default; value accrues by design.* An organization that runs the flow without managing it will therefore collect one hundred percent of the cost and an unknown fraction of the value. This unknown value fraction is not a gap in its spreadsheets. It is a gap in command and control. The purpose of AI Operations Management, stated in one sentence, is to close that asymmetry: to keep the usage flow visible through the

flow of records and to force a countervailing value flow to answer for the cost flow at every boundary that matters.

04

The functions of the discipline

A discipline is not a sentiment about a flow. It is a set of recurring functions, performed on a schedule, owned by someone, and measured. Supply chain management became real when it resolved into functions a firm could staff: demand planning, procurement, inventory control, logistics, distribution. AI Operations Management resolves the same way. Five functions, each answering a question the organization currently cannot answer, together constitute the practice of the discipline.

The first function is sourcing. Some function must decide which AI capacity the organization acquires, from which providers, under which terms, and for which work. The governing principle is that this is an economic decision rather than a technical ranking. The most capable model is not automatically the correct selection, because capability and economic suitability are distinct properties: a model can lead every benchmark and still fail a given task's requirements once total cost, latency, reliability needs, and the value actually at stake are brought inside the evaluation. Sourcing therefore matches work to capacity on economic terms, exactly as procurement matches components to suppliers rather than simply buying the finest grade of everything. The question it answers: Are we paying for AI capabilities we do not need or starving work that needs more?

The second function is planning and budgeting. AI budgets today are largely access budgets: they provision for seats and subscriptions and treat the consumption itself as beyond planning. A usage-aware budget does what a demand plan does for goods. It forecasts consumption by workflow, sets expectations against which actual usage can be compared, and defines in advance what growth is anticipated and what growth is anomalous, so that cost governance operates as control rather than as a quarterly surprise. Scaled AI deployment makes this a requirement, not a refinement, because exposure at scale moves faster than any retrospective review.

The question it answers: What do we expect this flow to consume next period, and how will we know when it deviates?

The third function is metering and attribution, and it is the discipline's central nervous system. Metering keeps the record flow running: every AI usage event is measured, recorded, and consolidated where the organization, not only the provider, can see it. Attribution converts those records into assignment. Cost that is assigned to a workflow, a team, or a decision has an owner. Cost that floats in an aggregate invoice has none, and assigning it requires an attribution basis chosen deliberately, because responsibility follows attribution. This is the function on which the other four depend, since sourcing cannot be evaluated, budgets cannot be compared, and value cannot be netted against cost without records that say what ran and for whom. The question it answers: Who consumed what, in service of which work?

The fourth function is allocation and routing. AI capacity is finite, on the provider's side and within the organization's own budget, and finite capacity under real demand must be allocated. Unmanaged allocation is the norm now, and it has created unmitigated chaos for AI strategy at most companies. Managed allocation becomes policy: which work is routed to expensive capacity and which to cheap, which usage carries priority when limits bind, which requests are throttled, queued, or declined. Providers already run their half of this function through tiers, limits, and priority pricing, which is one more reason the buyer's half cannot remain paralyzed. The question it answers: When the flow is constrained, who decides what runs first, and by what rule?

The fifth function is value-boundary management, and it is the one with no precedent in the software era, because software's economics never demanded it. Someone must draw the boundaries inside which AI's value is claimed to occur, and then hold the netting of realized value against fully loaded cost inside each one. A boundary is simply a fence drawn around a specific piece of the business, one workflow, one team, one quarter, inside which the question can be answered honestly: What did the flow cost here, and what did it return? Without an owned boundary, value claims float free of any test, productivity stories substitute for returns, and the organization runs the flow indefinitely without ever learning whether it pays. That is where we are right now with most Applied AI. The function does not need to prove value everywhere at once. It needs to ensure that

somewhere specific, someone is accountable for the answer. The question it answers: Where, exactly, does this AI flow pay for itself, and who is on the hook for knowing?

Five functions, one all-encompassing flow. Sourcing feeds it, planning anticipates it, metering and attribution see it, allocation disciplines it, and the value boundary makes it answer for itself. An organization performing all five is practicing AI Operations Management, whatever it calls the department. An organization performing none of them is not avoiding the discipline. It is abdicating one of the basic tenets of all business, cost control. We have already seen this play out. In 2025, MIT's State of AI in Business report found that 95 percent of enterprise generative AI pilots produced no measurable effect on the P&L, and S&P Global found the share of companies abandoning an AI initiative more than doubled in a single year. The spend arrived on schedule. The value was never made to answer for it.

05

What this is not

New disciplines are obligated to state their borders, because the territory near them is occupied and the occupants are doing legitimate work.

AI Operations Management is not AIOps, despite the collision of syllables. AIOps is the application of machine intelligence to the running of IT infrastructure: anomaly detection, incident prediction, automated remediation. In that field, AI is the instrument and infrastructure is the object. Here the relation is reversed. AI usage is the object, and the instruments are managerial: budgets, meters, attribution, allocation, boundaries. The two fields do not compete any more than a delivery truck competes with the logistics department.

It is not FinOps. FinOps is the nearest neighbor and the most natural ally, and its expansion into AI spend is real, rapid, and by its own account nearly universal. But its instruments are financial: visibility, allocation, forecasting, optimization. It sees the cost flow, and increasingly aspires to the value question, without owning the usage flow itself, deciding what work runs on what capacity under what priority, and without owning the value boundary, where realized value is made to answer

for accumulated cost. An organization can be excellent at FinOps and still collect all of the cost and an unknown fraction of the value. Cost visibility is the beginning of the discipline. It is not the discipline.

It is not MLOps or its descendants. MLOps governs the artifact: the model built, versioned, deployed, and monitored, the pipeline that carries it from training to production. This work is essential, and it is engineering. AI Operations Management begins where the artifact is already running and asks the questions engineering does not: what is flowing through it, at what cost, attributed to whom, and to what realized end. One discipline keeps the asset sound. The other manages the flow that the asset exists to serve.

Nor, finally, is it AI governance in the regulatory sense, the necessary work of risk, compliance, and responsible use. That discipline asks whether the organization should be doing what it is doing with AI. This one asks whether the organization knows what it is doing at all: what runs, what it costs, and what it returns. The answers to the second set of questions turn out to be the evidentiary foundation for the first, since certification of conduct rests on exactly the records, attribution, and boundaries this discipline exists to produce. But the disciplines are distinct, and this document claims only one of them.

Each of these fields governs an instrument, a budget line, an artifact, or a risk. None of them owns the flow. That vacancy is the entire case for a new discipline with a new name.

06

The Monday question

The argument of this document compresses to four sentences. AI use is resource consumption, not software access, and at scale it becomes one of the largest operational flows an organization runs. That flow carries cost automatically but value only by design. What cannot be seen inside the management boundary cannot be economically controlled. Flows of this size have always summoned a discipline, and this one has not yet been answered.

Whether your organization needs the discipline is not a matter of persuasion. It is a matter of five questions, and they can be asked on Monday morning.

Are you paying for capability you do not need, or starving work that needs more? What do you expect the AI flow to consume next period, and how will you know when it deviates? Who consumed what last month, in service of which work? When capacity is constrained, who decides what runs first, and by what rule? And where, exactly, does this AI flow pay for itself, and who is on the hook for knowing?

An organization that can answer all five is already practicing AI Operations Management, whatever it calls the practice. An organization that can answer none of them is not waiting to adopt the discipline. It is running its business in the dark, accumulating increasing costs by default, and deferring the question of value to a later date that its own records will never be able to reconstruct.

The goods were flowing long before anyone called it a supply chain. The naming was the beginning of the control. AI usage is flowing now, through your workflows, at your expense, toward outcomes no one is measuring. This document proposes the name.

AI Operations Management. The discipline of the flow.

A note on method. This document is the public face of a broader research project on the economics of business AI adoption. Every claim it makes rests on a documented research base, maintained so that each conclusion can be traced back to the premises that carry it. This volume asserts nothing that base does not support. Readers who want the underlying research are invited to ask for it.