Technology Doesn’t Create Value. The Value Chain (and TOGAF) Already Told Us That.

AI just makes forgetting it dangerous.

For the last 18 months, one of the most common things I hear in consulting rooms is: “We need an AI strategy.” Fair. Necessary. And also… a little risky when it starts from the wrong premise. Because the most expensive mistake I’m seeing right now is subtle: treating AI like it *creates* value, instead of what it actually does, which is optimize how value is produced.

If that sounds like semantics, it’s not. It’s economics. It’s governance. And in the AI era, it’s the difference between innovation and scaling a failure at machine speed.

We solved this problem decades ago

Long before “cloud-first,” “DevOps,” or “AI-native” became default vocabulary, we already had a clear answer to a basic question: *Where does value come from?*

Michael Porter gave us the clean version: firms create value by transforming inputs into outputs customers are willing to pay for. Technology matters, but it was never the source of value. It was a supporting activity. Important. Powerful. Still subordinate to the value chain.

Then TOGAF (and enterprise architecture thinking more broadly) sharpened this into practical operating language:

- **Utility (fitness for purpose):** did we deliver the outcome customers and the business needed?
- **Warranty (fitness for use):** can leaders trust it under real conditions: reliability, availability, performance, capacity, resilience?

Different vocabulary. Same truth: **value is external, production is internal, and technology is subordinate.**

AI does not change that. It just punishes us faster when we forget it.

Where DevOps and automation drifted (and KPIs locked it in)

DevOps and automation started with the right intent: reduce friction, improve flow, increase reliability. The failure wasn’t technical. It was economic and governance-related.

Here’s the drift pattern I’ve seen over and over:

We replaced value with activity metrics… then optimized relentlessly against them.

So we celebrated deployment frequency, lead time for change, pipeline "green," automation coverage, cost per transaction. None of these are "bad." They're just incomplete when they become the scoreboard.

The result is predictable:

- We measured **motion instead of meaning**
- We optimized **local success instead of system health**
- We improved **efficiency instead of value**

And yes, teams optimized perfectly… for the wrong thing.

AI is repeating the same pattern, just faster

Every DevOps failure mode now has an AI equivalent.

Then we measured deployment count.
Now we measure models deployed.

Then we celebrated pipeline green.
Now we celebrate "model accuracy."

Then we pushed for faster releases.
Now we push for faster autonomous actions.

Same mistake, new costume: treating internal efficiency as a proxy for external value.

The difference is brutal: **AI collapses feedback loops.** Errors used to arrive at human speed. Now they arrive at machine speed. If you automate a fragile process, you didn't transform it. You just made the failure arrive sooner (and louder).

Why AI governance is not optional

AI systems don't just execute tasks. They influence decisions, shape customer outcomes, and expand legal, ethical, and financial exposure.

That's why governance has to come back to first principles:

- **Utility:** Are we optimizing the *right* outcome, end-to-end?
- **Warranty:** Is the system trustworthy in real production conditions, including resilience and recovery?

When governance focuses on activity (tokens processed, throughput, number of models deployed), it can make AI look successful while the value chain quietly degrades.

Or said more plainly: dashboards can lie. Outcomes and reliability don't.

Replace broken KPIs with value metrics (that leadership can actually use)

A lot of AI metrics today are instrumentation metrics. Useful for engineers, but not sufficient for governance. Leadership needs metrics that map to Utility and Warranty, plus economics and accountability.

Here are practical replacements that keep you honest:

- **Instead of "model accuracy," measure "outcome correctness."**
  Accuracy can be high and still wrong in business terms (wrong customer, wrong context, wrong incentives). Outcome correctness forces the uncomfortable question: "Did we actually produce the right decision, for the right reason, in the real workflow?"
- **Instead of "tokens processed," measure "cost per reliable outcome."**
  Tokens are motion. Outcomes are meaning. And "reliable" matters because the hidden cost isn't just inference, it's incidents, remediation, rework, and customer damage when the system behaves badly. Scale without unit economics is just faster waste.
- **Instead of "% AI-assisted decisions," measure "decisions with defined owners."**
  This is the difference between *human in the loop* and *human on the hook*. Probability scores don't own consequences. People do.
- **Instead of "automation rate," measure "error detection latency."**
  If you want speed, great. But speed demands earlier detection, tighter guardrails, and faster rollback. Failing faster is only a virtue if you detect failures faster.

One simple rule that keeps governance clean: **if a KPI can't be mapped to Utility or Warranty, it doesn't belong in the governance framework.**

Not all AI carries the same risk, so oversight should scale by agent class

Governance gets practical when you stop treating "AI" as one category and start treating it like a portfolio of agents with different blast radii.

A useful risk lens:

- **Informational agents** (search, summarization): low risk, indirect impact
- **Assistive agents** (recommendations, drafting): medium risk, shapes decisions
- **Decision-support agents** (credit scoring, fraud decisions, underwriting): high risk, critical outcomes
- **Autonomous agents** (trading, fulfillment, ops control): extreme risk, direct action on the value chain

Rule of thumb: **the closer the agent is to executing value-chain decisions, the stronger Utility and Warranty controls must be.**

Autonomy without accountability isn't innovation. It's unmanaged risk.

The final test (that applies to every AI project)

Before signing off on the next major AI deployment, leadership should ask:

If this system is wrong, which customer outcome is harmed, how quickly will we know, and who is accountable?

If any part of that answer is unclear, the organization hasn't built a strategy. It has built a liability.

AI belongs where all technology belongs: optimizing value production, not redefining value itself.

Summary:

Modern organizations frequently mistake technology as a source of value rather than a tool to optimize the existing Value Chain. Historical lessons from DevOps and Enterprise Architecture reveal that focusing on activity metrics instead of Utility and Warranty leads to hollow efficiency. Artificial Intelligence accelerates this risk by collapsing feedback loops, potentially scaling failures at machine speed if it lacks proper accountability and governance. To avoid these pitfalls, leaders must shift from tracking technical outputs to measuring meaningful business outcomes. True innovation requires categorizing AI by its risk profile and ensuring that human responsibility remains intact even as systems become more autonomous. Overall, the text argues that AI belongs subordinate to the Value Chain, serving as a stabilizer and optimizer for human-defined goals.

Next week:

The production model every company is running — whether it admits it or not.

We're not rejecting AI.
We're restoring the discipline it requires.