# A.

# The End of the Transaction Factory

*Why shared services cannot optimize its way into the AI era*

Jim Scully | TheoryA.ai | July 2026

*This paper applies Theory A, a management philosophy for organizing agency in an age of abundant intelligence, to shared services. It is part of the AI Native Shared Services Collection.*

> *An operating model is a machine for managing scarcity. When the scarcity inverts, the machine does not merely become suboptimal. It becomes the constraint.*

For thirty years, shared services ran on a single economic engine: the same work, performed by cheaper labor, somewhere else. That engine is failing. Not because it was badly built. Because AI has changed the scarcity it was built to manage.

Knowledge retrieval, drafting, case preparation, routing, translation, routine analysis, first-line decision support. These are becoming abundant. The price of routine cognitive work is collapsing toward the cost of computation. And with it collapses the arbitrage that funds the legacy model.

This is not a technology problem. It is an organizational design problem. And the response most shared services organizations are choosing, AI inside the legacy model, will not solve it.

## THE SCARCITY INVERSION

Shared services was a rational architecture for a specific operating environment. Expertise was scarce, uneven, trapped in individual experience and email threads. The design response was exactly right: consolidate common work, specialize labor, segment demand into tiers, standardize processes, automate the repetitive, concentrate capacity where scale and labor cost could be managed.

The model was a machine for managing scarce human intelligence. That scarcity is now gone.

Consider what has become abundant in the last three years: knowledge retrieval and policy search. Plain-language explanation and summarization. Translation. Document drafting. Case classification and routing. Workflow initiation. First-line eligibility checks and decision support. This is, almost line for line, the work of Tier 0 and Tier 1, and much of Tier 2.

And consider what has become scarce: context, trust, judgment, governance, legitimacy, data quality,

policy logic, service design. Every mechanism of the legacy model economizes on the thing that is no longer scarce. None of them manages the things that now are.

Tiers ration expert attention. But expert attention embedded in policy logic no longer needs rationing. Process standardization protects consistency from unskilled variation. But consistency can now live in standardized constraints while paths personalize. Deflection pushed demand to cheaper channels. But resolution is now cheaper than deflection.

> *When the scarcity changes, the model does not merely underperform. It actively prevents the thing it claims to enable.*

---

THE FORK

Every shared services leader now faces a choice. The two paths use the same technology, which is why they are so easily confused in strategy decks and board updates.

The first path: AI inside the legacy model. Copilots in the towers, a chatbot on the portal, automation in the back office. The same operating model, run faster and cheaper. This path is measurable in familiar currency. It requires no structural change. It delivers real gains for several quarters.

The second path: redesign from the ground up. Not “which of our tasks can AI do?” but “if abundant intelligence had existed when shared services was invented, what would we have built?”

The first path is the natural choice. It is also a trap.

| | AI for efficiency | AI-native redesign |
|---|---|---|
| **Starting question** | Which of our tasks can AI do? | What must be true, and what system makes it true? |
| **Unit of accountability** | The ticket, the SLA, the handle time | The service outcome and its truth state |
| **Service structure** | Tiers retained. AI becomes smarter Tier 0 content. | Autonomy and risk segmentation. One interaction resolves what tiers passed upward. |
| **Consistency mechanism** | Standardized process paths | Standardized constraints. Personalized paths. Governance as architecture. |
| **View of demand** | Workload to deflect to cheaper channels | Signal to interpret and demand to prevent at the source |
| **Trajectory** | Early gains, then plateau | Compounding capability |

Read down the efficiency column and one pattern emerges: every row keeps the legacy model's assumptions intact and inserts AI beneath them. The AI-native column replaces the assumptions themselves. No sequence of efficiency gains adds up to that. Optimization of the old model is not a path to the new one.

---

THREE CLOCKS RUNNING

The efficiency strategy feels safe because it changes nothing structural. It is not safe. Three clocks are already running, and the efficiency path accelerates all of them.

**THE ARBITRAGE CLOCK.**
Arbitrage is a depleting asset. Every advance in AI shrinks the labor base the arbitrage is priced on. CFOs and clients will reprice accordingly, benchmark by benchmark, contract by contract. A shared services organization whose value proposition is labor cost gap multiplied by volume is optimizing a declining asset. Its own vendors, racing to automate the same work, will capture the difference first.

**THE TALENT CLOCK.**

A transformation framed as headcount reduction tells the workforce that AI's purpose is to replace them. The best people, the ones with judgment, the future experts, leave first. Meanwhile AI absorbs the routine casework that once trained novices into experts. If nothing replaces that apprenticeship, the pipeline of judgment quietly dies. The legacy model destroys the very scarcity the new model runs on, precisely when it matters most.

**THE DISINTERMEDIATION CLOCK.**

If shared services does not provide the enterprise's shared intelligence infrastructure, every function and business unit will build its own. Employees will route around a front door that only deflects, toward shadow AI channels with no governance. The enterprise ends up with fragmented intelligence delivery, inconsistent policy answers at machine speed, and no institution accountable for any of it. Shared services exists because fragmented service delivery was ungovernable. Fragmented intelligence delivery is worse.

> *Choosing efficiency is not preserving optionality. It is spending it.*

---

WHAT THE ALTERNATIVE ACTUALLY IS

## A coherent replacement exists. It can be stated in three substitutions.

| | Legacy model | AI-native model |
|---|---|---|
| **Accountability** | SLAs and service catalogs | Service outcomes with truth states. Conditions either hold or they do not. |
| **Structure** | Tiered support: Tier 0 through 3 | Autonomy and risk segmentation. Work classified by how much autonomy the system may exercise given the risk of the case. |
| **Consistency** | Standardized process paths | Standardized constraints. The boundaries are standardized. The path personalizes. Governance becomes architecture, not friction. |

Around these substitutions sits a different view of what shared services is for: not centralizing transaction execution, but providing the shared intelligence, policy logic, orchestration, governance,

and human expertise that let work happen safely and consistently anywhere in the enterprise.

The roles change accordingly. Case processors become service stewards, supervising a human-AI system, measured on escalation quality and judgment. Centers convert from labor-concentration sites to expertise, governance, and learning hubs. Agent operations becomes a formal discipline: agents designed, tested, monitored, tuned, governed, and retired like the workforce they are.

---

## WHY SHARED SERVICES, AND WHY FIRST

Every enterprise serious about AI eventually confronts the same question: where can a fundamentally different operating model be proven without betting the company?

Innovation labs fail this test because they own nothing real. Enterprise-wide transformation fails it because the blast radius is unbounded. Shared services passes it uniquely.

Shared services already has customers: employees, managers, functions, business units. It already has products: service lines, journeys, data services. It already has demand flows: cases, transactions, inquiries, escalations. It already has governance: service councils, controls, cost models, operating rhythms. And it carries production accountability of the most unforgiving kind. Payroll must run. Employees must onboard. Records must survive audit.

It is bounded enough to govern, real enough to prove anything, and measurable enough to convince a skeptical CFO.

The business case extends beyond shared services itself. Inside the unit, the returns are outcome achievement, waste elimination, and response latency. Beyond the unit, the return is option value: a proven AI-native shared services function delivers validated roles, governance patterns, agent operations, measurement frames, and management practices that de-risk the transformation of everything else.

> *The highest return on AI-native shared services is not the transactions it automates. It is the enterprise it teaches to operate differently.*

---

THE CLOSING PROVOCATION

AI does not make shared services obsolete. It makes the traditional rationale for shared services obsolete.

The model built to manage scarce labor cannot manage abundant intelligence. No volume of technology deployed inside it will change that. The structure, not the software, sets the ceiling.

Three clocks are running. Organizational readiness takes longer to build than AI capability. And the organizations that wait for the technology to mature before redesigning the vessel will discover, when the capability arrives, that they have nowhere to put it.

Theory A argues that organizations were designed around constraints that AI removes. Shared services is the clearest proof of that argument in the

enterprise, and the most powerful place to begin designing differently.

The question facing every shared services leader today is not whether to adopt AI. Everyone will adopt AI. The question is which side of the fork the adoption serves.

**One side optimizes what already exists. The other asks what shared services should have been all along, if it had been designed for a world of abundant intelligence from the start.**

Only one of those questions has an interesting answer.

---