

Beyond Placement: Why AI Fails Without Back-End Logic Reform

Ping Xu, GFI Flow Intelligence

In a previous PA Times article, I argued that placing AI at the front end of public systems—screening, classifying, or rejecting applications—does not improve decision quality. It amplifies error and pushes it upward. The principle was straightforward: the higher the ambiguity, the greater the need for human judgment.

That leaves a critical question: if AI should not act as the first decision-maker, where should it go?

The natural answer is the back end. But this answer is incomplete. Moving AI to the back end does not solve the problem if the back end itself is built on the wrong logic.

The Unseen Bottleneck: Specificity Logic

Most government back-end systems operate on what can be described as specificity logic. Each service is treated as a unique case. Each has its own workflow, its own integrations, and often its own customized implementation.

Under this logic, the marginal cost of adding a new service does not decline. The 100th service costs roughly the same as the first. Complexity accumulates, and integration overhead grows faster than linearly.

The result is not just inefficiency. It is systemic fragility. As services scale, failure rates increase, integration costs escalate, and agencies become locked into systems that are difficult to modify without significant reinvestment. What appears as digital transformation success at launch often becomes operational paralysis within a few years.

This is why many digital transformation efforts degrade over time. The system was never designed to learn. It was designed to repeat.

What AI Actually Enables: A Different Logic

AI—particularly large language models and intelligent agents—introduces a different possibility. Not simply faster processing, but a shift in system logic.

Instead of treating each service as unique, systems can be designed around shared patterns. Identity verification becomes a reusable module. Address changes become intents rather than forms. Eligibility checks become standard processes rather than custom-built workflows.

This shift—from specificity to universalization—changes the economic structure of the system.

Under universalization, marginal cost declines. The 100th service can be delivered at a fraction of the cost of the first. Complexity no longer compounds; it stabilizes or even decreases as the system scales.

This is not incremental improvement. It is a structural change in how systems are built.

Where AI Belongs in the Back End

When systems are built on universalization logic, AI plays three distinct roles:

1. Abstraction

AI identifies recurring patterns across services and helps define reusable modules such as identity verification, eligibility determination, and payment processing.

2. Intent Recognition

Instead of navigating forms, users express needs in natural language. AI maps diverse expressions to a finite set of intents, enabling consistent processing across services.

3. Dynamic Orchestration

AI assembles execution pathways at runtime, coordinating across modules and agencies without requiring fixed integrations between systems.

In these roles, AI operates in environments with lower ambiguity. It supports decision-making rather than replacing it. And it enables systems that improve with scale rather than degrade under it.

The Procurement Question No One Is Asking

Most government procurement processes are structured to validate delivery, not scalability.

Vendors are asked whether they can deliver service X—but not whether the system architecture can absorb service 50 without exponential cost. This omission is not technical. It is structural. And it explains why governments repeatedly procure systems that cannot evolve.

A more relevant question is: how does the marginal cost of adding services change over time?

This can be measured.

Two indicators are particularly useful:

- **Implicit Marginal Cost Reduction (IMCR):** the percentage decline in marginal cost as the number of services increases.
- **Generalization Level (GL Level):** a maturity scale indicating whether systems rely on shared modules, intent recognition, and dynamic orchestration.

These metrics shift evaluation from static delivery to system behavior over time.

A Procurement Correction

This gap can be addressed at the procurement level.

Governments do not need new technology mandates. They need new evaluation criteria. Requiring vendors to demonstrate declining marginal cost and measurable generalization levels would immediately change system design incentives.

Without this shift, AI adoption risks reinforcing structural inefficiencies rather than resolving them.

Automation applied to the wrong architecture does not fix the problem. It accelerates it.

Conclusion

The discussion around AI in government has focused heavily on capability—what AI can do. Less attention has been given to placement and underlying system logic.

AI does not belong at the front end of high-ambiguity decisions. That principle is increasingly clear.

But it also does not belong in a back end built on specificity logic. In such systems, AI simply automates fragmentation.

The real challenge is not deploying AI. It is ensuring that the system it operates within is structured to benefit from it.

Universalization provides that structure.

Without it, AI will not transform government systems. It will make their limitations more efficient.