

AI-Economy-Architecture

GUDIYA Cognitome | V2 Reconstruction | September 2026

An architecture for commitments and their consequences

An agent accepts a sales order. Other agents reserve inventory, request transport, and schedule delivery. Each digital action changes what someone else must fulfill. The order may travel through several enterprises before the customer receives it, while its consequences compete with thousands of other commitments for the same resources.

GUDIYA proposes an architecture that makes these relationships visible and governable. It connects a network of digital actants with a network of physical resources, follows the effects of initiated work, and uses observed and forecast constraints to inform admission and ongoing stabilization. The purpose is to keep the creation of obligations compatible with the ecosystem's ability to fulfill them.

The five figures develop this architecture progressively:

- structure,
- differences in speed,
- continuity across a workflow,
- propagation of disturbances, and
- resource saturation.

They are conceptual illustrations. Their hexagonal geometry is a visual language for connectivity, not a requirement that real networks have regular cells or equal numbers of neighbors.

Two connected layers

Figure 1 presents the upper digital layer and the lower physical resource layer. Digital actants include agents, software services, APIs, and the digital interfaces through which people or embodied systems participate. Some workflows remain digital only. Others depend on resources represented below, such as production equipment, inventory, warehouses, vehicles, and energy systems.

The lower layer represents resources and their relationships: what they can provide, what they require, and what limits their operation. A warehouse depends on transport and labor; a factory depends on materials and power. Connections between layers carry instructions and commitments downward, and observations, constraints, acknowledgments, and settlement evidence upward.



Figure 1. The AI-Economy-Field. Digital actants and physical resource actants form two connected networks, with selected couplings between them.

A physical resource and its digital representation must have an explicit identity relationship. Otherwise the same capacity could be counted twice. The architecture also distinguishes ownership from dependence: an enterprise may rely on a partner's vehicles or warehouse without owning them. Its effect graph can cross those organizational boundaries even when visibility into the partner's operations is limited.

The picture simplifies a richer system. Compute, storage, and network bandwidth are finite resources too. Resource accounting therefore applies to digital activity as well as physical fulfillment. The two-layer view helps explain the economy without implying that only the lower layer can become constrained.

Different speeds within one economy

Figure 2 introduces a frequent mismatch: digital actants can create and revise commitments much faster than the associated physical work can settle. A dispatch request may be created almost instantly, while picking, loading, travel, and receiving require elapsed time. More requests do not automatically increase the capacity available to complete them.



Figure 2. The upper layer operates rapidly while the lower layer fulfills work more slowly. The contrast depicts a common relationship, not a universal timing rule.

Neither layer has a single clock. Some physical control actions are fast; some digital processes take hours. The Grid should therefore associate settlement expectations with operations and obligations rather than assign one permanent speed to an entire enterprise. Relevant measures include deadlines, available service windows, expected duration, uncertainty, and the freshness of reported state.

Following a Gudiya-Effect

A Gudiya-Effect comprises an initiated action and its attributable consequences across a declared scope, followed through to verified resolution of its settlement obligations. Figure 3 follows one such effect between digital actants and into physical fulfillment. The Gudiya-cookie links records of the same effect across these transitions.



Figure 3. A Gudiya-cookie preserves continuity as a Gudiya-Effect traverses digital actants and physical resources. Repeated GE-01 markers identify the same effect at different points in its lifecycle.

The illustrated sequence is deliberately simple. An operational effect is usually a graph: it can branch, merge, retry, and create subordinate obligations. A single transport operation may serve several effects. The records must preserve those relationships without counting the shared resource consumption several times. The cookie is the continuity connector; associated records hold the graph, evidence, and accountable owners.

Explicit couplings include declared calls, reservations, orders, and handoffs. Implicit couplings include shared capacity or state through which one actant changes another's conditions without directly messaging it. Two purchasing agents may never interact, yet their orders may compete for the same loading bay. Inferred relationships need confidence and provenance; a shared identifier alone does not establish causality.

Settlement means that the outcomes and obligations specified for the effect have been verified as resolved, or that remaining responsibility has been explicitly transferred and accepted. A timeout or completed API call is insufficient evidence. The declared scope prevents "complete settlement" from becoming an impossible demand to track every consequence indefinitely.

Disturbances can travel through both layers

Figure 4 introduces a perturbation at an upper-layer actant and shows potential propagation through digital and physical dependencies. A revised schedule can cause rerouting, additional reservations, and new congestion. The response at one resource can then change the decisions made by digital actants elsewhere.



Figure 4. A perturbation propagates through upper-layer couplings and crosses into the lower resource network. The highlighted routes depict possible transmission, not inevitable collapse.

A perturbation is a disturbance and can arise inside or outside the chosen system boundary. Instability describes the resulting behavior: deviations may grow, harmful oscillations may persist, or the system may leave acceptable operating bounds. An external disruption can trigger internal amplification, while an internal retry loop can create a disturbance without an external event. The boundary must be stated before calling an origin endogenous or exogenous.

Potential instability patterns include repeated overshoot and correction, runaway retries, congestion feedback, resource starvation, and cascading fulfillment failures. A large disturbance need not produce instability if buffers and controls absorb it. Conversely, small disturbances can matter when delays, gain, and shared dependencies reinforce their consequences.

Saturation limits what can be promised next

Figure 5 returns to the tracked effect and adds saturated resources. The red admission markers concern new commitments. Existing accepted work still needs fulfillment or accountable recovery. Closing admission does not cancel what has already been promised.




Figure 5. Saturated physical resources refuse additional commitments while existing obligations continue toward settlement. The Gudiya-cookie retains continuity through the accepted workflow.

Constraint	What can become exhausted
Throughput	The rate at which work can be completed.
Stock	Available inventory, stored energy, or other consumables.
Concurrency	Simultaneously usable machines, slots, bays, or execution capacity.
Time	Capacity available before an obligation’s deadline.
Recovery margin	Reserve needed to handle faults, exceptions, or disruption.

Capacity is specific to a resource, service, location, and time window. A facility may be unable to accept one kind of order while retaining capacity for another. The Grid must represent those distinctions, including reservations and uncertainty, instead of reducing every resource to one universal percentage-full indicator.

How the Grid governs proposed and active work

An intent-group is a proposed execution plan submitted for approval. It should identify the desired outcome, participating actants or selection rules, expected obligations, timing, resource needs, and settlement conditions. An approved intent-group is a plan admitted under stated constraints. Adaptive changes that materially alter its dependencies or demands require reassessment.

The Grid evaluates new work against existing obligations and forecast capacity, with the relevant resource owners confirming what they can support. Approval may be conditional, time-limited, or staged. Uncertain future demand should remain distinct from accepted commitments, and forecasts must not silently become reservations. Distributed reservation records need a way to prevent different participants from spending the same capacity.

The Gudiya-Effect graph follows work through its lifecycle. A dynamic capacity and commitment graph shows where multiple effects contend for shared resources over time. DDMC can use both views to estimate remaining margins; neither graph can be assumed complete. Missing or stale capacity information must remain visible as uncertainty.

Discipline	Role in this architecture
Admittance Control	Establish participation conditions and admit new commitments within authorized limits.
DDMC	Distributed Dynamic Margin Computation estimates capacity, deadline, and recovery margins under current and anticipated demand.
DLPC	Distributed Dynamic Leverage Points Calculation evaluates where a bounded intervention may best reduce propagation or contention.
Cognitive Rx	Apply approved pacing, resequencing, duplicate suppression, reconciliation, or coordinated substitution.
Circuit Breaking	Interrupt selected harmful actions or couplings while preserving essential functions and responsibility for unfinished work.

The operational loop is continuous: observe state and obligations, predict pressure, evaluate plans, reserve or pace work, verify outcomes, and revise the forecasts. A plan acceptable at admission may become unsupportable after a capacity loss. The Grid may then defer new work, reschedule active work, or suspend selected actions within its authority. Any resulting residual obligations still require an owner.

Buffers need backing and limits

Flywheel services can buffer bursts of proposed actions, suppress retries, and coordinate orderly draining of work. At settlement, their counterpart needs usable reserve capacity, scheduling slack, or agreed alternatives. A digital coordinator cannot supply a missing truck or manufacture inventory. Every buffer needs a limit, an acceptable waiting interval, and an escalation action before it is exhausted.

Coordination can itself amplify disruption if all participants follow the same erroneous forecast or reroute simultaneously. Local safeguards, staged interventions, feedback freshness, and fallback behavior are therefore part of the architecture. The objective is accountable stabilization under uncertainty, not a claim of perfect prediction.

Enterprises can participate gradually

The two-layer representation supports different depths of participation. An enterprise may expose digital actants first, add one constrained resource later, and expand only when the operational value justifies the added integration. Another enterprise may retain that narrower scope indefinitely. Grid services and assurances must match the visibility and authority actually provided.

Participation	Information shared	Potential Grid support
Digital	Actant identities, interaction telemetry, and declared obligations.	Pacing, retry control, and observable digital dependency management.
Acceptance envelope	A bounded quantity or service capacity for a defined window, with conditions and expiry.	Admission coordination without disclosure of internal operational details.
Selected resources	Chosen resource capacities, reservations, dependencies, and settlement evidence.	Targeted bottleneck prediction and balancing.
Broader resources	Wider connected resources and explicitly authorized intervention options.	More extensive coordination of fulfillment and recovery.

For example, an enterprise might report that it can accept another 200 orders in a specified delivery window without revealing staffing, equipment, or internal utilization. It remains responsible for computing and honoring that acceptance envelope. The Grid needs clear rules for confirmation, revision, and expiry so the envelope cannot be consumed repeatedly by different requesters.

Three decisions remain separate: what the enterprise reveals, what the Grid may infer or coordinate, and what the Grid may change. Sharing telemetry does not automatically authorize operational rescheduling. Detailed commercial information should remain private where possible, while affected partners receive only the constraints and commitments necessary for coordination.

An enterprise with substantial physical dependencies can still choose limited visibility. Undisclosed capacity remains unknown; it cannot be certified as sufficient. Participation flexibility must not allow an enterprise to bypass the conditions of a resource owner whose capacity it requests. Resource owners retain the ability to confirm or refuse new obligations.

Two structural layers and two lifecycle responsibilities

Digital and physical layers describe the structure of the economy. Commitment and settlement describe responsibilities throughout a Gudiya-Effect's lifecycle. They should not be equated. Digital-only work can require settlement, and a physical process can generate further commitments. Both responsibilities operate across the two structural layers.

What would establish the value of this proposal

A bounded pilot should compare existing coordination with Grid-assisted operation under equivalent demand and disturbances. Track overdue obligations, capacity conflicts, duplicate commitments, recovery time, and unnecessary rejection. Test stale telemetry, partial participation, and coordinator loss. The architecture earns its value when it helps participants make supportable commitments and complete settlement without exporting hidden burdens to their neighbors.

Two Ways to Participate in the GUDIYA Grid

An enterprise can begin with Digital Agents Only. By connecting its software agents to the Grid, it could gain help coordinating their actions, suppressing duplicate requests, controlling excessive retries, and tracking unfinished obligations. Consider an online furniture retailer: the Grid could help prevent its sales, purchasing, and delivery agents from repeatedly changing the same order or issuing conflicting instructions. These benefits require visibility into digital activity, without necessarily exposing the retailer's detailed physical operations.

Adding Resource Agents extends that visibility to fulfillment. A resource agent represents something with finite capacity—inventory, a warehouse, a production line, or a delivery fleet—and reports its relevant availability and constraints. The Grid could then recognize that the retailer's agents are accepting orders faster than its warehouse can dispatch them, or that furniture is available but tomorrow's trucks are fully booked. Within agreed authority, it could help pace new commitments, propose achievable delivery windows, or coordinate alternative capacity before promises become overdue obligations.

Digital participation helps stabilize what the software is doing. Resource participation additionally helps reconcile what the software promises with what the operation can fulfill. Both support settlement tracking, but resource participation can provide stronger evidence about the conditions affecting actual fulfillment. Neither guarantees successful outcomes; the usefulness of coordination depends on timely information, reliable resource commitments, and clearly assigned responsibility.

Adoption can be gradual. An enterprise might connect its digital agents first, then add only its most constrained resource. It could also share a limited acceptance envelope—"We can accept another 200 orders for tomorrow"—without disclosing staffing, equipment, or internal utilization. Sharing information does not automatically authorize the Grid to change operations. Each enterprise defines its participation scope, while the Grid's assurances remain bounded by what it can observe and influence.

What type of Instabilities could originate/propagate through the AI-economy

Trigger origin is relative to the participating digital and physical network represented in the GUDIYA Grid. “Both” means either an internal or external trigger is possible. The layer of origin identifies where the instability mechanism typically begins.

Type of instability	Typical layer of origin	Possible trigger origin	How it can originate	How it can propagate
Retry amplification	Layer 1 — Digital	Both	Internal delays or external service failures prompt repeated requests after ambiguous responses.	Additional requests deepen congestion, provoke further retries, and may duplicate physical obligations.
Recursive task proliferation	Layer 1 — Digital	Endogenous	Agents delegate, spawn subtasks, or invoke one another without effective bounds.	Expanding branches consume compute and create additional downstream commitments.
Decision oscillation	Layer 1 — Digital	Both	Agents repeatedly reverse decisions in response to one another, stale observations, or changing external conditions.	Suppliers, schedules, reservations, and physical operations undergo repeated rearrangement.
Conflicting control actions	Layer 1 — Digital	Endogenous	Multiple agents independently control the same process or shared state.	Their actions interfere, triggering corrections, rework, and incompatible physical instructions.
Correlated decision amplification	Layer 1 — Digital	Both	Many agents respond similarly to a shared model, forecast, or external announcement.	Simultaneous purchasing, cancellation, or rerouting concentrates demand on shared resources.
Prediction–reality divergence	Layer 1 — Digital	Both	Modeling errors or environmental changes invalidate predictions that continue to guide action.	Plans create unfulfillable obligations; corrective responses can amplify the mismatch.
Digital congestion feedback	Layer 1 — Digital	Both	Internally generated activity or external demand overloads processing, causing delays and timeouts.	Dependencies stall; retries increase load, while stale instructions disrupt physical execution.
Fulfillment backlog growth	Layer 2 — Physical resources	Both	Overcommitment or capacity loss leaves work arriving faster than it can be completed.	Delays spread to dependent operations; agents may respond with retries, duplicate bookings, or overordering.
Inventory depletion feedback	Layer 2 — Physical	Both	Internal allocation problems or external supply disruption provoke hoarding, larger	Apparent demand rises, exhausting alternative suppliers and spreading

Type of instability	Typical layer of origin	Possible trigger origin	How it can originate	How it can propagate
	resources		orders, or duplicate reservations.	shortages.
Resource deadlock or starvation	Either layer	Endogenous	Work holds resources while waiting for others, or allocation repeatedly excludes particular obligations.	Settlement stalls, reservations remain occupied, and dependent workflows accumulate.
Overload-induced capacity deterioration	Layer 2 — Physical resources	Both	Internal scheduling or an external demand surge produces loading that increases errors, fatigue, breakdowns, or rework.	Effective capacity declines further, worsening overload and shifting pressure to neighboring resources.
Recovery-resource exhaustion	Either layer	Both	Ordinary work, internal failures, or external disruptions consume reserves needed for recovery.	Local disruptions persist and spread because alternatives, repair capacity, or fallback services are unavailable.
Commitment–settlement divergence	Cross-layer interaction	Both	Digital commitments accumulate faster than fulfillment can absorb them, through overcommitment or lost capacity.	Outstanding obligations grow across suppliers, warehouses, transport, and customers—even with lightly loaded computers.
Delayed-feedback oscillation	Cross-layer interaction	Both	Internal latency or external communications disruption causes agents to issue corrections before earlier actions take effect.	Overshoot triggers reversal, producing repeated swings in production, dispatch, or allocation.
Premature settlement and release	Cross-layer interaction	Endogenous	Completion rules mistake an acknowledgment for verified fulfillment and release reservations too early.	Capacity is promised again while still occupied, creating overlapping obligations and corrective activity.
Rerouting cascade	Cross-layer interaction	Both	Multiple workflows respond to an internal bottleneck or external disruption by choosing the same alternative.	The alternative saturates, prompting further rounds of rerouting across the network.
Cancellation and compensation amplification	Cross-layer interaction	Both	Internal replanning or external cancellations withdraw commitments after execution has begun.	Returns, refunds, recalls, replacement orders, and reconciliation create additional work and obligations.
Stabilization-induced instability	Grid intervention across either layer	Endogenous	Aggressive throttling, simultaneous restarts, or poorly targeted circuit breaking disrupts useful activity.	Intervention shifts overload elsewhere, creates oscillations, or obstructs settlement and recovery.

Interpretation: This is a proposed classification of possible mechanisms, not a finding that every occurrence becomes unstable. Saturation can remain stable under controlled admission. An external disturbance can trigger internal amplification, and visible damage may appear far from its origin. For actual incidents, record the evidenced trigger origin—or Undetermined—separately from the propagation mechanism.