

117-311-D Grid-Node Capacity

A Multidimensional Foundation for GUDIYA Stabilization Capacity

Stability Engineering Institute | Version 1.0 | 6 October 2026

Capacity to remain responsible

How many actants can a GUDIYA node handle? The question appears straightforward until we ask what handling means. A node may register an identity, receive a heartbeat and store a record without being able to follow the actant's changing relationships, maintain orientation, detect an emerging instability or intervene before a consequential commitment. The capacity that matters is the capacity to discharge the stabilization responsibility that attachment creates.

SEI proposes that Grid capacity should be described through a multidimensional capacity profile and an operating envelope. The profile states available resources and demonstrated service capabilities. The envelope states which combinations of workload, coupling, deadlines and operating conditions those capabilities can support. An enterprise, regional node, hyperscaler region or national Grid therefore needs more than an advertised maximum number of agents.

This is a foundational engineering proposal. Its dimensions, measurement conventions and acceptance tests require development and validation. It builds on established capacity-planning practice while extending the planning boundary to orientation, adaptive interactions, runtime emergence and the full Gudiya-cookie lifecycle. HCAS here means a Complex Adaptive System operating at machine speed.

The airport that cannot grow outward

As an example, the challenge of capacity for an airport, begins with a constrained airport: demand grows even when additional land is unavailable. Its most useful observation is that an airport contains several capacities. Runways, taxiways, gates, passenger processing and surface access must work together. A larger terminal cannot resolve a saturated runway, and a runway improvement can expose a new constraint in baggage handling or ground access. [1]

This also shows why the unit being counted matters. A larger aircraft can carry more passengers without adding another aircraft movement, yet increases the work arriving at the gate and terminal. In GUDIYA, a more productive cognition-cookie may likewise carry more subagents, rendezvous, reservations and consequences. Fewer cognition-cookies need not mean less stabilization work.

Aviation already distinguishes benchmark capacity from capacity available under prevailing conditions. FAA material describes operational arrival capacity as changing with configuration, weather, traffic mix and other constraints. [2] The GUDIYA inference is that a node's declared capacity must similarly identify its workload and operating regime. An ideal-condition figure should not silently become a promise during a partition or a burst of interacting failures.

Three quantities that must remain distinct

1. Cognitive capability concerns what an actant can understand or accomplish.
2. Execution capacity concerns the resources available to perform its work.

3. Stabilization capacity concerns the Grid's ability to observe, interpret, coordinate and constrain that work within its mandate. Increasing model intelligence or inference throughput can increase the burden on stabilization without increasing the capacity to provide it.

What the airport analogy contributes

Airport constraint	GUDIYA counterpart	Capacity lesson
Runway movements and separation	Consequential commitments and intervention spacing	The permissible action rate can bind before compute is exhausted.
Gates and turnaround time	Active cognition-cognition-cookies and retained obligations	Longer residence consumes capacity even without more arrivals.
Taxiways and intersecting routes	Rendezvous, dependencies and shared services	Interaction geometry matters alongside participant count.
Terminal and baggage processing	Observation, analysis and durable evidence	Faster ingress can overwhelm the services that follow it.
Weather and equipment condition	Uncertainty, failures and reference degradation	The same infrastructure supports different operating capacities.
Neighboring airports and access links	Federated nodes and downstream resources	Local expansion must fit in the wider network.

These are functional analogies, not claims that cognitive dynamics obey aerodynamics. Airports themselves are sociotechnical systems with adaptive participants. The useful transfer is the discipline of examining coupled constraints, workload mix and changing operating conditions rather than reducing capacity to a single piece of infrastructure.

A working definition

Grid stabilization capacity is the set of workloads a specified node or federation can sustain, for a stated duration and operating regime, while satisfying its observation, orientation, intervention, settlement and recovery obligations under human-defined policy. "Sustain" includes continued useful operation and a controlled response to the disturbances covered by the capacity claim.

Counts answer different questions.

- Registered actants are an inventory;
- concurrently active actants are an operating population.
- Cognition-cookie arrivals per second are a flow;
- unsettled cognition-cookies are a stock.
- Detected emergencies are tracked objects;
- discovery work also includes candidate patterns that never become recognized objects.

Each count needs its own definition, population boundary and time interval.

The same enterprise can own actants in all three layers of the SI-Economy. Layer membership follows the function being performed, not a permanent claim that every agent of one company has the same consequence. Capacity planning must attribute the work actually performed, including the work that a single cookie induces in several layers.

The supply behind the services

The underlying supply vector includes processor time, accelerator time where used, working memory, durable storage, read and write throughput, network bandwidth, power and operational staffing. Shared identity services, clocks, enforcement interfaces and external dependencies can impose additional constraints. Their capacities are measured in their own units; bytes, processor-seconds and staff-hours cannot be added into an undifferentiated total.

Twelve dimensions for a first capacity profile

Google's SRE guidance makes a relevant engineering point: requests per second can be a misleading capacity proxy because request costs vary. It recommends accounting for the resources consumed. [3] GUDIYA should retain that discipline and add the responsibility-specific constraints that determine whether adequate stabilization is being delivered.

The following dimensions describe functional capacities that a Grid operator should characterize. They are a Version 1.0 planning vocabulary, not twelve independent machines or universal constants. Each capability consumes underlying resources, and several may compete for the same processor, storage system or network link. Limits must be measured together under a declared workload mix.

Dimension	What must be measured	What exhaustion means
1. Attachment and identity	Registered and active identities; admission, renewal and revocation rates; identity state size.	Attachment or authority changes become delayed, inconsistent or unverifiable.
2. Cookie lifecycle state	Concurrent active and unsettled cognition-cookies; residence times; bytes and obligations retained per cookie.	The node cannot reliably retain the context and commitments it has accepted.
3. Rendezvous management	Formation and reconfiguration rates; active groups, group size and coupling churn; delinking work of two or more actants together.	Participants cannot safely establish, change or conclude shared activity.
4. Observation transport	Events and bytes per second, burst size, timestamps, loss, ordering and maximum evidence age.	The node's evidence becomes incomplete or too old for its responsibility.
5. Orientation and estimation	CCS state updates and AOI deliveries per second; inference cost, reference freshness and uncertainty.	Actants or controllers act against stale or inadequate operating references.
6. Emergence discovery and tracking	Candidate-analysis budget; object birth rate; live objects, memberships, update cost and retirement lag.	Consequential collective patterns cannot be discovered or followed adequately.

Inventory is not workload

A million dormant identities may be easier to maintain than a small group repeatedly changing its authority, membership and resource commitments. Conversely, a large dormant population can still impose renewal, security and recovery costs. A useful profile separates resting inventory from active processing and specifies what happens when many dormant participants become active together.

A Gudiya-cookie is the accountable undertaking and its associated lifecycle record. Its cost depends on the work and obligations represented, not merely the identifier's storage size. A cookie with many children, cross-node rendezvous and unresolved external effects is a different workload from a short local task. Capacity must include closure work even after the originating actant stops running.

Completing the capacity profile

Dimension	What must be measured	What exhaustion means
7. Control and enforcement	Interventions per second; reachable targets; command-to-effect delay; verification and simultaneous action limits.	Detection remains possible while effective correction becomes too slow or unavailable.
8. Commitment and settlement	Allowed commitment rate; confirmation throughput; unresolved outcomes, reserved resources and reconciliation work.	Obligations accumulate or commitments outrun the ability to verify and settle them.
9. Federation and handoff	Cross-node traffic; state-transfer size; handoffs and coordination rounds; consistency and jurisdiction constraints.	Work cannot safely move or coordinate despite spare capacity elsewhere.
10. Evidence and replay	Durable ingest, retention and query volume; lineage links; replay and audit-processing throughput.	Accountability, diagnosis and recovery lose required evidence.
11. Recovery and continuity	Reserved processing and communication; failover load; resynchronization backlog; power and outage endurance.	The node functions in normal conditions but cannot survive a stated failure.
12. Human assurance and operations	Staffed concurrent incidents; escalation response; approval and review queues; maintenance coverage.	Exceptional situations exceed the human authority and operational support available.

A vector is the beginning of the model

Write the available resource profile as C and the resource demand induced by a workload as D . A necessary resource check is that demand in each dimension remains within its available budget. If a reserve has already been removed from the available budget, it must not be subtracted again. Shared consumption must be accounted for once at the underlying resource, even when it supports several functional dimensions.

$$D_i(\text{workload}, \text{regime}) \leq C_i(\text{regime}), \text{ for every resource } i$$

This comparison is necessary but insufficient. A node can remain below every average resource limit while missing an intervention deadline, losing a critical coupling from its field of observation or operating outside its authority. A capacity claim therefore requires a feasible operating region: combinations that satisfy resource, timing, evidence-quality, enforceability and policy constraints together.

This region need not be rectangular or fixed. Increasing discovery effort may reduce the processing available for routine estimation. Adding a remote controller may add compute while increasing communication delay. A list of separately measured maxima cannot be treated as a promise that all maxima can be achieved simultaneously.

The geometry and velocity of demand

A workload descriptor should include active populations, task arrival rates, task durations, communication fan-out, coupling density, recursion depth, topology churn, common dependencies and the rate of consequential commitment. It should also identify consequence severity, reversibility and the time available for correction. These descriptors explain demand; they are not interchangeable physical resource units.

Two populations with identical headcounts can impose very different loads. Independent local tasks may require limited coordination. The same actants sharing one scarce resource may generate contention, retries, correlated decisions and repeated envelope changes. A new subagent can add both its own workload and new obligations for the group that spawned it.

Potential pairwise relationships grow rapidly with population, but actual work does not automatically require examining every possible pair. Many interactions are sparse; some important effects are mediated by a shared resource rather than a direct message. Multi-actant rendezvous also creates group constraints that pairwise counts alone miss. Benchmarks must vary the actual interaction structure and measure the cost of discovering relevant unknown relationships.

Emergence creates a new capacity obligation

Once GUDIYA represents a consequential emergent pattern as a First-Class Grid Object, that object requires an identity, scope, evidence, lineage, trajectory estimate and review or retirement conditions. A tightly synchronized retry wave may need more frequent updates than a slow trend. Some objects overlap, merge or split. Tracking cost therefore depends on membership, update rate and analysis complexity, not just the number of named emergences.

Discovery must be budgeted separately from tracking. Counting only already-detected objects would make the Grid appear lightly loaded precisely when its discovery machinery is overwhelmed. A proposal for the research program is a protected discovery budget, measured against test patterns of known provenance. False positives, missed patterns and investigation cost must all appear in the results. Treating an object as a useful representation does not prove a causal explanation of it.

Time places a limit on usable capacity

An overloaded stabilization service may continue processing every event eventually and still fail its purpose. Evidence must arrive, be interpreted and produce an effective intervention while a corrective opportunity remains. Queueing delay, communication delay, decision time and actuator response belong in one end-to-end timing budget, with an explicit uncertainty margin.

$$\textit{Observation} + \textit{queueing} + \textit{analysis} + \textit{decision} + \textit{actuation} \leq \textit{deadline} - \textit{margin}$$

The deadline comes from the protected process and operating regime. A percentile latency target describes a service level, not a hard safety guarantee. Hard deadlines require a defensible bound or an independent protective mechanism. Nyquist is relevant only where sampled-signal assumptions apply; neither twice an actant's token rate nor a faster observer by itself proves adequate observation or control.

When evidence deteriorates, the Grid may need to narrow permitted activity, reduce commitment velocity or invoke local protective controls. Slower inference is not the only available control. The design question is which change restores the required margin while preserving useful operation.

An illustrative node at its limits

Consider a hypothetical node supporting 20,000 active actants. It accepts 500 cognition-cookies per second, with an average time to accountable closure of 12 seconds, and each cookie induces eight rendezvous-management operations. It tracks 400 emergent objects at 25 updates per second each. The figures below are invented planning inputs, not GUDIYA measurements or recommended universal frequencies.

For this example, available budgets are assumed to have already allowed for emergency reserves. Functional limits are assumed to have been tested jointly under this particular mix. Identity, compute, power, enforcement and other constraints remain mandatory even though the table displays only six dimensions. An operating node would need evidence for every such assumption.

Capacity measure	Baseline demand	Illustrative available budget
Active actants	20,000	25,000
Active or unsettled cognition-cookies	$500/s \times 12\ s = 6,000$	10,000
Observation events	$40,000 + 10,000 + 10,000 = 60,000/s$	100,000/s
Rendezvous operations	$500/s \times 8 = 4,000/s$	6,000/s
Emergence tracking updates	$400 \times 25/s = 10,000/s$	20,000/s
Cookie closure throughput	500/s	800/s

Observation demand assumes two actant events per second, twenty additional lifecycle events per arriving cookie, and one observation event per emergence update. These categories are disjoint for this arithmetic. Processing an event and analyzing it still consume separate service resources. More updates cannot manufacture new independent evidence.

The baseline cognition-cookie inventory uses Little's law: under suitable long-run flow conditions, average inventory equals the accepted arrival rate multiplied by average residence time. [4] If residence becomes 30 seconds while that rate remains sustainable, average retained cognition-cookies rise to 15,000, exceeding the illustrative limit. This average relation is not a transient backlog forecast or a bound on peaks.

Now consider separate stress cases. If tracked emergencies rise to 1,200, tracking demand becomes 30,000 updates per second, beyond its budget, although observation traffic reaches only 80,000 events per second and actant count is unchanged. If each cookie instead requires sixteen rendezvous operations, demand becomes 8,000 operations per second, beyond that service's limit. Neither failure is visible in an actant-only rating.

If closure capacity temporarily falls to 300 cognition-cookies per second while 500 continue entering, unresolved inventory grows by 200 per second: another 12,000 in one minute, absent other exits. Delaying settlement may sometimes protect an external system, but the retained obligations need capacity. Admissions, reservations and completion policy must be adjusted together; a protective delay cannot become an unbounded queue.



Capacity follows the cookie across the three layers

Layer 1, the Agentic Economy, introduces planning, delegation, model execution, discovery and interaction demand. Layer 2, Digital Services, introduces service queues, databases, APIs, transaction processing and digital commitments. Layer 3, Resource Services, introduces constrained resources such as energy, equipment, inventory, transport and physical operating schedules. Some resources are reservable long before their actual use. A cognition-cookie can cross these boundaries repeatedly.

A purchasing undertaking illustrates the interaction. A planning actant compares alternatives, inventory and payment services reserve a transaction, a warehouse reserves handling capacity, and a changed delivery condition sends the undertaking back for replanning. One customer intention can generate multiple loops through all three layers. Capacity estimation must follow these paths and their retry or compensation branches, not count only the first request.

Grid processing capacity and downstream economic capacity are different constraints. A node may be able to authorize a million reservation requests while the warehouse can honor only a small fraction. A free Layer-1 attachment also has a nonzero infrastructure cost. The proposed layered funding model determines who pays; it does not determine how much observation and control can safely be omitted.

Approval is also a resource commitment

During Intent-Group Approval, the Grid should specify a recommended stability envelope and assess the capacity needed to uphold it. The approved undertaking should declare expected participation, permitted delegation, dependencies, workload class, timing requirements, resource reservations and completion conditions. Estimates can be revised within authority as adaptive behavior changes the undertaking.

Provisional, read-only and action attachments have different demand profiles. Read-only participation may avoid external commitment while still consuming identity, telemetry, reference and analysis capacity. The node should reserve enough capacity for its existing obligations before accepting new ones. If a novel undertaking has uncertain cost, a bounded trial can establish evidence before broader admission.

Approach, link, operate, delink and settle

Rendezvous Management needs capacity for the whole interaction lifecycle. Establishing a group consumes authorization and orientation work. Active collaboration consumes observation and coordination. Membership change requires reevaluation. Delinking can require release of shared reservations, isolation of unfinished effects and reorientation of remaining participants. These operations do not disappear because setup was successful.

Closure releases capacity only when the relevant obligation has actually ended. A terminated agent may leave children, a reservation, an uncertain payment outcome or a resource still in motion. A node must retain and reconcile those obligations, while permitting verified completed portions to retire according to policy. This is the Grid equivalent of gate turnaround: reducing avoidable residence can increase useful capacity without adding hardware.

Reducing residence must not mean declaring success prematurely. Higher closure throughput is beneficial only when completion evidence and residual obligations remain correct. Resource release and evidence retention may have different lifetimes and therefore require separate planning.

The scale changes the planning boundary

Planning scale	Primary responsibility	What must be included
Enterprise node	Support the enterprise's mixed workloads and business commitments.	Application dependencies, local GRMs, peak calendars, recovery objectives and any shared external services.
Regional node	Coordinate several enterprises or local nodes with shared constraints.	Cross-enterprise coupling, correlated demand, domain expertise, inter-node traffic and regional continuity.
Hyperscaler region	Provide tenant isolation and large-scale execution and stabilization services.	Burst spawning, shared control planes, network partitions, noisy neighbors, inference contention and failure domains.
National Grid	Maintain a wider stability view and coordinate cross-sector responsibility.	Regional summaries, systemic dependencies, inter-grid boundaries, escalation capacity and national contingency arrangements.

These are planning scopes, not four necessarily consecutive levels in one hierarchy. A hyperscaler region can span many enterprises and interact with several sector or geographic authorities. Ownership of compute does not automatically confer authority over every consequence. The capacity model must identify the responsible controller and jurisdiction as well as the host.

Federation does not make capacities simply additive

Two nodes with spare processing cannot necessarily absorb each other's workloads. A handoff may require state transfer, access to the same enforcement points, compatible policy, data-location permission and enough time to establish orientation. Shared power, networks, identity systems or external services can make their nominal redundancy fail together.

A national rating must therefore include the network of dependencies and feasible placement of work. Spare capacity in a distant node is useful only if the relevant undertaking can use it within its timing and authority

constraints. Capacity assigned as reserve for one failure scenario must not be simultaneously promised to several regions that may fail together.

National coordination need not ingest every raw event or issue every local control command. Local GRMs should handle fast, bounded stabilization, while regional and national functions exchange the evidence needed to recognize wider coupling and shared pressure. However, aggregation can conceal local structure or correlated behavior. The adequacy of summaries must be tested against the systemic patterns the receiving controller is expected to detect.

The Grid also consumes its own capacity

Observation, reference distribution, replication, health checks, discovery and auditing consume resources. Dependencies between Grid controllers add coordination traffic and failure modes. Planning must include this overhead and protect the control path from ordinary actant demand. Shared visibility must be sufficient for each responsibility, with acknowledged uncertainty. Connectivity alone does not establish complete national awareness; inadequate evidence must narrow the corresponding authority to act.

A practical planning method

1. define the service promise. Specify the node's responsible domain, workload classes, observation coverage, orientation freshness, intervention deadlines, consequence limits and closure obligations. State which failures and degraded regimes the promise covers. Capacity has little meaning until the responsibility against which it will be judged is explicit.
2. build a workload catalogue. Describe representative cognition-cookies from approval through closure, including their subagents, multi-actant rendezvous, external commitments and retained histories. Separate inventory from flow, background activity from bursts, and mandatory control work from optional analysis. Include discovery of new emergence as a standing workload rather than an exception that is assumed to be free.
3. measure a service-demand model. For each workload class, profile processor time, memory residence, communication volume, storage and operator intervention. Measure coupling and coordination overhead under mixed workloads rather than adding isolated benchmarks. Fit a provisional demand model, retain uncertainty ranges and revise it when the workload or implementation changes.
4. examine time horizons. Annual growth and procurement determine installed capacity. Seasonal peaks, tenant launches and sector schedules determine placement and reservations. Runtime admission and pacing govern the next seconds or minutes. The same plan must accommodate abrupt topology changes and the time needed to provision, attach, orient and warm a replacement node.
5. model operating regimes and correlated stress. Test ordinary peaks, prolonged residence, burst spawning, rapid group reconfiguration, sudden emergence discovery, external-service slowdown and loss of a node or link. Vary observation quality and actuator delay as well as hardware utilization. A simulation explores consequences of its assumptions; it does not validate dynamics that its model cannot represent.
6. establish reserves and allocation policy. Protect observation and intervention for already-admitted work, preserve capacity for safe completion or transition, and bound retries and new admissions. Reserve percentages should follow measured scenarios, uncertainty and service consequences. There is no justified universal rule that every Grid should run at a particular utilization level.
7. validate the complete loop and publish a qualified capacity profile. Move from replay and shadow observation to controlled workloads, staged stress and limited operational authority. Record both the boundary at which obligations fail and the cost of conservative restrictions. Increase admitted scope only when evidence supports the revised profile.

The capacity declaration

A node's declaration should state hardware and software versions, workload mix, active and retained populations, relevant interaction structure, measurement interval, sustained and burst limits, latency and evidence-quality targets, operating regime, protected reserves, covered failures and expiry or review conditions. An actant count can appear as a derived result for that profile; it must not stand alone as a universal specification.

Human staffing belongs in the declaration when exceptional cases depend on it. Humans can authorize bounded policies in advance while machine-speed loops execute within that mandate. Nevertheless, a node that generates exceptions faster than its operators can resolve them may accumulate blocked commitments or continue under uncertainty. Automation does not erase that capacity constraint.

Growth without pretending capacity is unlimited

The airport example offers several strategies beyond expansion [1]

- reduce avoidable occupancy,
- improve sequencing,
- reorganize conflicting paths,
- move suitable work away from constrained locations,
- and add replacement capacity before retiring old facilities.

For GUDIYA, these suggest

- shortening verified cookie residence,
- improving rendezvous scheduling,
- partitioning unnecessary coupling,
- moving nonurgent replay off the control path and
- staging upgrades with spare capacity.

Each improvement must be judged end to end. Batching can reduce overhead but increase waiting time and the size of correlated bursts. Compression can reduce bandwidth but conceal a signal needed for detection. Partitioning can isolate failures while making cross-boundary coupling harder to see. Larger cognition-cookies can perform more useful work while enlarging the consequences and recovery burden of one undertaking.

More hardware can help only where hardware is the binding constraint. Other limits may require

- a better estimator,
- additional enforcement coverage,
- clearer authority,
- a narrower mandate,
- an improved workload architecture
- or reduced commitment velocity.

Capacity engineering should expose which constraint is active before an operator purchases the wrong form of expansion.

Installed, demonstrated and currently available

Installed capacity is the inventory of resources. Demonstrated capacity is the workload region supported by a stated body of tests. Currently available capacity is what the surviving configuration can offer now, after reservations, maintenance, failures and uncertainty are taken into account. These quantities should never be collapsed into one dashboard number.

Islanded operation is a distinct capacity profile. Loss of the wider Grid can reduce visible dependencies and available evidence even when local processors are idle. Local mandates and disjoint resource budgets may preserve useful work, but remote reservations and global consistency cannot be assumed. The node needs a declared duration of autonomous operation and a capacity budget for reconciliation when connectivity returns.

Recovery can be more demanding than normal operation:

- reconnecting actants refresh their references,
- deferred cognition-cookies seek settlement,
- histories replay and services retry.
- Admission during recovery should be paced against the reunited system's actual capacity.
- Spare resources must exist long enough to drain the backlog while serving the essential ongoing load.

Avoid making overload self-reinforcing

Google's overload discussion describes how retries and connection bursts can intensify demand and contribute to cascading failures. [3] In GUDIYA, the concern extends to adaptive reactions: actants may create helpers, seek alternative services or repeatedly replan when delayed. Those reactions should be represented in the stress workload rather than treated as unrelated traffic.

The local GRM should use bounded queues, reservation expiry, workload-aware admission and explicit backpressure, with policy for safe continuation or interruption. It must avoid creating a flood of notifications that overwhelms its own control path. Reducing optional work can preserve essential stabilization, but suppressing critical evidence to make utilization appear healthy is not a capacity improvement.

What SEI should measure next

The first research result should be an experimental capacity surface, not a universal actants-per-node number. Hold the node implementation fixed and vary population, cookie arrival rate, residence time, coupling structure, group size, topology churn, emergence birth rate, update cost and commitment deadlines. Repeat under failures and recovery. The resulting surface shows which combinations remain within the stated obligations.

Report useful cognition-cookies closed correctly alongside active obligations, oldest evidence age, end-to-end intervention delay, deadline misses, envelope excursions, discovery precision and recall against labeled test cases, false interventions and time to recover. Measure subgroup and worst-case behavior rather than relying only on fleet averages. Count completed interventions by verified effect, not merely by commands sent.


The Stability Unit also needs a precise meaning. If it denotes a deployed stabilization component, it should carry a qualified multidimensional profile. If SBU denotes a normalized unit for metering or exchange, its reference workload, quality level and accounting interval must be specified and calibrated. A commercial summary unit can aid comparison and billing, but it cannot establish that a particular mix of obligations is supportable.

SEI should investigate how much emergence-related demand can be forecast from leading indicators, which observations are sufficient for a delegated scope, how reserves should change with uncertainty, and when partitioning improves capacity without hiding consequential coupling. These are empirical questions. Evidence may show that some GUDIYA functions are cheaper, costlier or less effective than this first architecture assumes.

The capacity principle

A GUDIYA Grid node should accept only the work for which it can maintain the required orientation, observation, intervention and accountable closure under the declared conditions. The Grid's capacity is consequently a responsibility bounded by evidence. It is a vector of available capabilities, constrained further by the interactions and deadlines of the work entrusted to it.

The airport analogy leads to a practical objective: increase useful activity by identifying and relieving the actual constraint while preserving the operating margin. For GUDIYA, success means enabling more constructive cognition without admitting more responsibility than the stabilization infrastructure can fulfill.



Sources

- [1] The Challenge of Capacity, Mentour-Now talk transcript. The airport bottleneck, occupancy, sequencing and expansion examples inspired this chapter. Advertising segments and unverified airport rankings or numerical claims were not used.
- [2] Federal Aviation Administration, CDM International Online Workshop, Lesson 5: Capacity, Delay, Weather and Contingencies. Supports the distinction between benchmark and operational capacity. <https://tfmlearning.faa.gov/tfm-training/atfm-basics/cdm-t1-lesson5b.html>
- [3] Google, Site Reliability Engineering, Chapter 21: Handling Overload. Supports resource-based accounting and the need to manage retry, connection and overload effects. <https://sre.google/sre-book/handling-overload/>
- [4] MIT, Urban Operations Research, Section 4.4: Some Important Relationships in Queueing Theory. Supports the average flow–inventory relationship and its conditions. https://web.mit.edu/urban_or_book/www/book/chapter4/4.4.html
- Web sources accessed 6 October 2026.