

117-300-D GUDIYA-AOI

Actant Orientation Interface with Gridwide visibility

SEI Concept Proposal • Version 2.0 • 3 October 2026

A pilot must make decisions from inside an aircraft while understanding conditions that extend beyond it. A helmet-mounted display can place relevant flight and sensor information within the pilot's field of view. For an F-35, this includes horizon, airspeed and altitude information wherever the pilot looks. The display makes instrument-derived references available at the point of decision; it does not create those references or guarantee correct interpretation. [1]

The Stability Engineering Institute proposes an analogous function for adaptive actants: GUDIYA-AOI, the Actant Orientation Interface. SEI proposes this as a built-in capability for every Grid-attached actant. It would give each actant a current, scoped and attributable view of its undertaking, operating limits, relevant shared conditions and outstanding obligations. For a software agent, this is principally a machine-readable interface. A human-facing display may accompany it, but a literal wearable device is unnecessary.

The Grid should inform the actant it observes

GUDIYA's wider observation is useful for intervention, but it can also improve decisions before intervention becomes necessary. An actant may know its objective and local tool results while lacking evidence of congestion, competing reservations or commitments created by its delegates. Returning relevant Grid context can help the actant adapt within its mandate instead of repeatedly discovering constraints through failed actions.

Gridwide visibility means that relevant wider observations inform the interface. It does not mean perfect knowledge or unrestricted disclosure of every participant's activity. The Grid estimates conditions within its coverage; GUDIYA-AOI delivers a scoped view to the actant. Authorized governance tools can receive corresponding views directly from the same trusted Grid context service. This chapter develops that proposed architecture and its engineering requirements.

Orientation in the cognitive space

Orientation means understanding the undertaking's current state relative to its purpose, authorized envelope, dependencies and unresolved effects. This need not imply that cognition has literal altitude, compass direction or a universal physical coordinate system. Domain-specific references can provide useful orientation before any broader theory of cognitive measurement is established.

The pilot analogy also has a limit. Aircraft orientation draws on mature physical measurements. Cognitive operating variables, their meanings and their relationship to stability require definition and validation. The FAA's guidance on spatial disorientation illustrates the importance of reliable instrument references when bodily sensations are misleading. It does not establish a neurological equivalent in an AI system. [2]

A reference for adaptive operation

In SEI's terminology, HCAS is a Complex Adaptive System operating at machine speed. Plans, collaborators and resource demands can change faster than a static briefing remains useful. The AOI should make consequential changes visible while preserving continuity of identity and obligations. It supports informed adaptation; it cannot eliminate uncertainty or replace enforceable control.

GUDIYA-AOI : Actant's Dashboard

A car's dashboard gives its driver selected information needed to operate: speed, fuel remaining, engine temperature and warnings. The driver does not need to inspect every component continuously. Instruments bring relevant conditions into view while the driver concentrates on the journey.

GUDIYA-AOI proposes an equivalent dashboard for an actant operating in the cognitive space. For a software agent, this dashboard is primarily a structured, machine-readable stream of references rather than a visual display.

Driver's dashboard	Actant's dashboard
Speedometer	Current activity rate, commitment rate and permitted pace
Fuel gauge	Remaining compute, resource and delegated expenditure budgets
Engine temperature	Signs of operational stress, congestion or accumulating unresolved work
Navigation display	Approved objective, lifecycle stage and current dependencies
Distance and collision warnings	Developing interaction conflicts and proximity to stability-envelope boundaries
Warning lights	Expired authority, unreliable evidence, lost Grid connection or required intervention
Instrument status	Freshness, provenance and uncertainty of the displayed information

The actant's dashboard must reflect both local conditions and relevant ecosystem conditions. An actant may be functioning correctly while many other actants are placing unsustainable demands on the same service or resource. Its own successful execution is therefore insufficient evidence that continued acceleration is appropriate.

The dashboard should also return evidence of consequences: what completed, what remains pending, and what obligations survive a change of plan. When reliable information is unavailable, it must explicitly indicate uncertainty rather than display an apparently healthy state.

The dashboard informs decisions; enforcement mechanisms apply binding limits. GUDIYA-AOI connects the actant's local intelligence to the wider operating picture, helping it remain oriented as conditions change.

The actant dashboard

A car's dashboard helps its driver operate through selected indicators. GUDIYA-AOI proposes an equivalent dashboard for an actant: structured data delivered at the point of decision, with a visual display available for human oversight.

Driver reference	Actant reference	Operational question
Speedometer	Activity and commitment rates; permitted pace.	Is further acceleration appropriate?
Fuel gauge	Remaining compute, resource and expenditure budgets.	What capacity remains within the mandate?
Temperature gauge	Congestion, unresolved work and operational stress.	Are demands accumulating faster than resolution?
Navigation display	Objective, cookie stage and active dependencies.	Where does the next action belong?
Proximity warning	Interaction conflicts and envelope margins.	Is a boundary or coupling hazard approaching?
Warning lights	Revoked authority, stale evidence or lost connection.	What needs attention or intervention?
Instrument validity	Evidence source, age, coverage and uncertainty.	Is the operating picture reliable enough?

A healthy local view may conceal wider stress

An actant may be meeting its objective while many other actants are exhausting the same service or resource. Its successful tool calls do not establish that continued expansion is sustainable. The dashboard should therefore combine relevant local facts with Grid-derived assessments of shared conditions, including changing margins and unresolved collective exposure.

Information should support a decision

Each indicator needs an operational definition and source. Begin with measurable rates, reservations, delays and obligations. Estimated boundaries should state their uncertainty and assumptions; undefined cognitive speed or universal stability scores would offer false precision.

The dashboard should return consequences as well as intentions: completed effects, pending actions and obligations that survive replanning. When evidence is missing, it must indicate unknown or stale state. The actant should not interpret silence as permission or a healthy operating condition.

Orientation and enforceable control

The dashboard informs decisions. Binding controls remain independently enforceable. The Grid may issue an advisory that helps the actant choose among permitted alternatives, or an authorized restriction that an execution system must apply. Clear message types preserve this distinction and allow useful adaptation within human-defined policy.

What the interface should convey

The following mapping translates functions rather than physical quantities. Each deployment must choose variables appropriate to its consequences, operating horizon and available evidence.

Pilot information	Proposed actant reference	Decision it supports
Attitude and horizon	Current state relative to the authorized envelope and its assumptions.	Is the undertaking still operating within supported bounds?
Position and heading	Cookie identity, lifecycle stage, objective and direction of relevant change.	Where does this action belong in the continuing undertaking?
Airspeed and operating margins	Rates of commitment, resource exposure, queue pressure and recovery margin.	Can another action be sustained now?
Nearby traffic and hazards	Relevant counterpart readiness, shared dependencies and emerging conflicts.	Will this interaction interfere with other activity?
Warnings and communications	Typed advisories, restrictions, rendezvous conditions and control commands.	What should change, and what must change?
Instrument validity	Source, timestamp, expiry, uncertainty and coverage limitations.	Is this reference fit for the pending decision?

A proposed orientation record

A record should identify its issuer, intended recipient, actant lineage and Gudiya-cookie. It should carry an observation time, issue time, sequence or version, validity period and the governing mandate and envelope versions. Values need defined meanings and units or categories. Missing information should remain explicitly unknown rather than being silently interpreted as normal.

The payload should distinguish observed facts, estimated conditions, forecasts, recommendations and binding instructions. It should state the affected scope, remaining reservations or obligations and any required acknowledgement. A forecast with uncertain assumptions must not look identical to a confirmed service state. A recommendation must not be mistaken for permission to expand authority.

Show what matters without exposing everything

The actant needs sufficient context for its decisions, not the entire Grid's telemetry. Shared capacity warnings can often be conveyed without revealing another participant's identity, private objective or detailed activity. Access controls and data minimization apply to the orientation service itself. Different actants may receive different views while referring to compatible definitions and versions.

The interface should prioritize consequences and changes. Repeated low-value messages consume attention, tokens and compute and can hide urgent information. A concise summary should retain access to supporting evidence where permitted, with explicit indications of what the summary omits.

Where the interface sits in GUDIYA

The AOI is a functional interface, not a requirement to create a separate product for every supporting responsibility. Existing runtime services may implement parts of it if they provide the required semantics, trust boundaries and response behavior.

Grid element	Relationship to the AOI
Intent-Group Approval	Establishes purpose, participants and delegated powers; binds approval to the Grid's recommended and authorized envelope.
Actant identity and lineage	Keeps orientation attributable across delegates, restarts, host changes and model replacement.
Grid-responsive attachment	Provides authenticated exchange of observations, orientation updates, acknowledgements and control outcomes.
GSE-Lang	Defines interoperable meanings for state, restrictions, validity, requests and responses.
Cognitive Coordinate System	Provides shared references for comparing relevant state and change; dimensions require operational definitions.
Rendezvous Management	Applies current authority and conditions to a proposed interaction, including readiness and reservations.
Gudiya-cookie and stability record	Preserve the undertaking's history, envelope versions and unresolved obligations through settlement.

A continuous exchange with distinct responsibilities

During restricted registration, the Grid supplies a test orientation record. Operational activation follows approval and verified responsiveness. The actant reports relevant state and proposed changes, while the Grid incorporates independent service and resource evidence where available. The Grid then updates its assessment and sends revised references. The actant can revise its plan or request a rendezvous using that context. Consequential execution still passes through the applicable authorization and enforcement mechanisms.

Receipt, interpretation and compliance are different events. A delivery acknowledgement establishes only that a message arrived. The runtime may need evidence that the current version is active, that a restriction has been enforced and that the expected effect occurred. Those acknowledgements should be represented separately.

Orientation is not authorization

A fresh view of available capacity does not itself reserve that capacity. Conditions can change between receiving the record and committing work. Rendezvous or another authoritative execution check must revalidate consequential assumptions and perform any required reservation. Orientation improves planning; it does not remove concurrency conflicts.

Similarly, the AOI should not rely on placing advisory text in a prompt as its only means of control. Binding restrictions must reach an enforcement point outside the actant's discretion. The actant may misunderstand or ignore a warning, yet the restriction must still apply where the architecture claims enforceability.

Intelligence needs continuing orientation

Reasoning capability does not guarantee awareness of the operating situation. An actant can generate a plausible plan while holding an obsolete picture of resource availability, another participant's readiness or its own outstanding commitments. In an adaptive environment, yesterday's accurate reference can become today's misleading premise. GUDIYA-AOI is intended to make that changing context available at the point of decision.

Aviation offers an engineering lesson: orientation requires reliable references as well as an operator capable of using them. Flight instruments provide references when unaided perception is unreliable; a helmet-mounted display can make relevant information accessible within the pilot's view. The corresponding GUDIYA proposal is to design orientation into the actant's operational architecture, rather than assume that intelligence alone supplies it.

The perception action and feedback relationship

The actant receives a scoped orientation record, interprets it against its mandate and proposes or performs an authorized action. Execution systems report what actually happened. The Grid combines those effects with wider observations, reassesses conditions and returns updated references. Consequences therefore inform the next decision, including consequences beyond the actant's immediate tool response.

This relationship is reciprocal. Actants contribute observations and discrepancies; the Grid contributes a wider assessment. Neither view is infallible. The interface must preserve the distinction between a requested action, an accepted request and a confirmed effect. Local reasoning and independent enforcement remain necessary throughout this exchange.

No GUDIYA AOI means no operational Grid attachment

SEI proposes that every actant participating in the GUDIYA-stabilized cognitive space must have a functioning GUDIYA-AOI and demonstrate Grid responsiveness. The requirement concerns capability and behavior, not the presence of a branded software package. Implementations may vary, provided they satisfy the applicable interface and control requirements.

This is a stabilizability condition within admission control. Identity checks establish who is participating; permission checks establish what that participant may do; responsiveness checks establish whether the actant and its execution environment can operate within the Grid's governed feedback relationship. Passing one class of check does not substitute for the others.

Restricted registration resolves the attachment dependency

Before operational attachment, a restricted registration channel permits identity verification, GSE-Lang capability negotiation and harmless responsiveness tests. Conditional Intent-Group Approval specifies the proposed undertaking and envelope, including required response times and degraded behavior. The actant receives operational authority only after the necessary tests and approval conditions are satisfied.

The test covers the actant, its AOI connection and its enforcement environment together. A responsive interface attached to an execution system that cannot apply restrictions is insufficient. The proposed policy is therefore: no GUDIYA-AOI, no operational attachment; no demonstrated Grid responsiveness, no activation.

The Grid Responsiveness Challenge

A CAPTCHA suggests a useful challenge–response analogy, but the Grid is not trying to distinguish a human from a machine. It is testing whether the registered participant can receive current references and respond within its declared operating requirements. The following challenges are proposed engineering tests, not an existing certification standard.

Harmless challenge	Required evidence	Capability assessed
Deliver a scoped record	Acknowledge the intended identity, envelope version and expiry.	Correct reference binding and interpretation.
Revise a test limit	Observe the new limit taking effect at the covered execution point.	Responsive enforcement rather than verbal agreement.
Pause a test task	Confirm covered work pauses within its specified deadline.	Timely control and observable effect.
Expire a reference	Observe entry into the agreed degraded mode.	Safe operation when orientation is lost.
Replay an older record	Confirm it cannot restore superseded authority.	Version discipline and replay resistance.
Restore current guidance	Confirm authorized resumption with obligations preserved.	Controlled recovery without duplicate effects.

Verify effects rather than assurances

An actant replying that it understands is insufficient evidence. The Grid should compare acknowledgements with independent execution telemetry where available. Test records should be fresh and bound to the participant, session and tested configuration, so an old successful response cannot simply be reused. Evidence should identify what was tested, the result, its validity and any untested scope.

A pause test must distinguish stopping new work from cancelling already accepted remote work. Where a remote commitment cannot be cancelled, the runtime must expose that residual obligation and follow the approved containment or completion procedure. A model that stops generating text has not thereby stopped every downstream effect.

Continued authority requires continued responsiveness

Successful admission demonstrates behavior under tested conditions. Adaptive plans, new delegates, changed permissions, runtime updates and migration can alter the relevant control relationship. Consequential changes should trigger reassessment and, where needed, renewed challenges before expanded authority is granted. Low-impact runtime checks should supplement actual operational evidence without destabilizing useful work.

If responsiveness is lost, the agreed policy may restrict new commitments, isolate affected scope or suspend authority. Necessary completion and reconciliation channels should remain available where policy permits. Disconnection does not erase the Gudiya-cookie's obligations or prove settlement. Reattachment requires current references and reconciliation of unresolved effects.

Shared context for authorized lifecycle tools

GUDIYA-AOI is the actant-facing interface to a broader Grid context service. The same service can supply purpose-specific views to authorized tools governing the undertaking. Each consumer must implement the agreed semantics and be authorized for the particular data and scope. Receiving observations does not confer authority to change policy.

Consumer	Relevant Grid context	Potential response
Actant through AOI	Objective, dependencies and envelope margins.	Revise a plan within the mandate.
Security system	Trust changes, lineage and suspicious interaction patterns.	Investigate or apply an authorized restriction.
Workflow orchestrator	Shared exposure, pending work and congestion.	Adjust concurrency, delegation or sequence.
Rendezvous controller	Counterpart readiness, authority and reservations.	Approve, defer or constrain a coupling.
Settlement controller	Outstanding obligations and verified effects.	Hold closure until conditions are satisfied.
Human oversight	Developing conditions and intervention outcomes.	Review evidence or revise the mandate.

Continuous cybersecurity after attachment

Identity can remain unchanged while tools, collaborators and exposure change. A trusted update can report that a service is now suspect, delegated authority has expired or a sequence of permitted actions warrants investigation. Detection requires supporting observation and analysis; the interface distributes the resulting context. Novel behavior alone is not evidence of malicious intent.

For example, an actant may retain configured access to a supplier API after credible evidence indicates that the endpoint is compromised. The Grid could distribute that evidence to the actant and a security policy controller. Under delegated authority, the controller could restrict further access and report whether enforcement took effect. Previously accepted remote work would still require reconciliation.

A proposed integration with NVIDIA Sentry

NVIDIA describes Sentry as correlating interactions, policy decisions and tool or data access, with DOCA connecting programmable BlueField protection to OpenShell policy. Its overview also describes continuing identity and delegated-authority checks. These functions suggest possible integration points; the overview does not establish a GUDIYA-AOI ingestion interface. [3]

SEI proposes a direct trusted Grid feed to Sentry or its policy controller, alongside the actant's AOI view. The security system should not depend on the actant forwarding its dashboard. An adapter would need to translate authenticated, scoped Grid updates into supported policy operations and return enforcement evidence. This is an integration proposal requiring validation, not a claim of existing product compatibility.

The Grid feedback loop

The AOI feed is a delivery path inside a feedback loop. The loop closes when observations of actual consequences return to the Grid and inform its next assessment. A delivered message, an acknowledged instruction and an effective intervention are distinct events.

Stage	Responsibility	Evidence carried forward
1 Observe	Actants, services and independent sensors report.	Activity, state changes and confirmed effects.
2 Assess	Grid relates local and ecosystem conditions.	Estimated state, couplings, margins and uncertainty.
3 Communicate	AOI and tool interfaces distribute scoped updates.	Versioned context, advice or authorized commands.
4 Respond	Actants adapt; execution systems enforce.	Actions taken and restrictions applied.
5 Verify	Grid compares outcomes with intended effects.	Residual exposure, discrepancies and revised assessment.

Fast local control with wider coordination

Immediate stabilization can remain local while the Grid supplies broader context and constraints. Validated, time-bounded updates can be applied locally without asking a remote authority to approve every step. Human policy defines delegated powers, operating limits and degraded modes. Grid learning must not silently expand those powers.

Timeliness depends on evidence age, observation coverage, communication, assessment, enforcement and effect-verification delays. The relevant recovery window determines whether the complete loop is adequate. Nyquist is a specific sampling result under defined signal assumptions; monitoring faster than an agent's nominal action rate does not by itself prove adequate observation or stable control.

A feed that can be trusted

A common contract should bind each update to its issuer, recipient, cookie, authority scope, observation time, version and expiry. It should identify whether the payload is an observation, estimate, forecast, advisory or binding instruction. Provenance and uncertainty help recipients avoid treating an inference as a confirmed state.

Policy precedence and permitted issuers must be agreed before conflicting updates arrive. Stale messages must not restore revoked authority. Out-of-order delivery, gaps, duplicate events and clock uncertainty require explicit handling. Authentication, integrity protection and independent enforcement protect the feed from becoming a channel for forged Grid instructions.

Feedback can also introduce instability

If many actants react identically to the same warning, their simultaneous corrections may create oscillation or move congestion elsewhere. The Grid should assess the effects of its own interventions. Where justified, bounded rates of change, hysteresis, staged release and differentiated instructions can reduce overreaction. These controls require domain-specific validation; a common feed alone does not stabilize an ecosystem.

An undertaking crossing the three SI Economy layers

Consider a procurement actant seeking a replacement component before a production deadline. Its Gudiya-cookie links planning in the Agentic layer, procurement and logistics calls in Digital Services, and inventory and transport commitments in the Resource layer. It can return to planning while earlier branches remain active. The Grid assesses cross-layer conditions; AOI delivers the relevant view. Planning can change while earlier commitments persist.

Event	Orientation made available	Bounded response
Initial plan	Objective, deadline, approved exposure and current evidence requirements.	Explore alternatives within the mandate.
Supplier delay	Service delay is confirmed; two inventory reservations remain open.	Account for existing holds before creating another branch.
Adaptive delegation	Delegates share one collective exposure budget; membership and reservations are current.	Replan without multiplying independent claims on the same budget.
Resource deterioration	Transport capacity estimate has changed; earlier availability is no longer reliable.	Seek a new rendezvous or authorized objective revision.
Loss of reference	A required source is stale; new commitments are restricted under the agreed degraded mode.	Continue only permitted bounded work and reconcile outstanding effects.
Completion request	A successful delivery is recorded, but an unused reservation lacks release confirmation.	Keep settlement pending until release or accepted transfer is verified.

Cross layer orientation and a concrete message

Agentic intent, service execution and resource consequences remain linked. A resource constraint can require service throttling and agentic replanning. A human-readable update might say: Your procurement authority remains valid. The shared capacity estimate has deteriorated. Two delegated reservations remain active. New reservations are temporarily restricted by the current envelope. Resolve an existing hold or request reassessment before expanding exposure.

The actual exchange should carry structured fields identifying the issuer, effective version, scope and expiry. The explanation helps the actant choose a useful next action; the restriction is separately enforceable. If an older record later arrives, it must not silently restore superseded authority.

AI Vertigo as an operational metaphor

SEI can use AI Vertigo to describe a mismatch between an actant's local interpretation and the wider operating situation. Examples include acting on an obsolete plan, assuming a delegate has stopped, or treating an accepted request as a completed effect. This is a proposed operational metaphor, not a clinical diagnosis or proof that AI has human sensations.

AOI can expose mismatches through current references and contradiction notices. A shared false Grid view can also worsen them. Independent checks and explicit uncertainty remain essential.

Requirements for a trustworthy interface

Orientation data is a consequential input. Messages need authenticated origin, integrity protection, recipient and scope binding, and defenses against replay or unauthorized substitution. An actant should not be able to promote ordinary tool content into a Grid instruction. Access to the interface must not expose authority or private information beyond the recipient's mandate.

Freshness and degraded operation

Update timing should follow the dynamics and consequences of the application. Evidence age, transport delay, clock uncertainty and processing delay all affect whether a reference remains useful. An update frequency alone does not demonstrate adequate observation or timely intervention. The full path to effective action must fit within the relevant recovery window.

Each deployment should define what happens when references conflict, expire or disappear. A conservative local mode may limit new commitments while completing essential accepted work. It should preserve ownership of residual obligations and provide a route to reconcile after reconnection. The correct response is not always immediate termination of every activity.

Avoid making the Grid a shared source of error

A common mistaken reference can produce correlated decisions across many actants. The Grid should report uncertainty, maintain provenance and allow discrepancies to be raised. Independent service evidence, local checks and controlled comparison of estimates can reduce shared blind spots. Learning from a discrepancy must not silently rewrite human policy or expand an actant's authority.

Coverage across the three SI Economy layers

Layer	Relevant orientation	Cross layer significance
Agentic	Intent, delegation, collaborators and collective commitments.	Parallel plans may create competing claims on shared capacity.
Digital Services	Availability, permissions, queues, pending requests and transaction status.	Accepted service requests may still await resource execution.
Resource	Compute, energy, inventory, reservations and physical limits.	Capacity changes may require service restrictions and agentic replanning.

The cookie can revisit any layer, branch across several layers concurrently and carry unresolved effects from earlier steps. Orientation must preserve that continuity. At each layer, the Grid should relate the immediate interaction, the actant's history and relevant collective patterns, while explicitly reporting gaps in observation.

Validating the integrated architecture

A bounded trial should compare otherwise similar agents with and without the AOI while retaining the same security and enforcement controls. Test stale records, contradictory evidence, delayed acknowledgements, dynamic delegation and migration. Include an actant that ignores advice, a Grid outage and a case where the Grid's estimate is wrong.

Measure harmful or unnecessary commitments, unresolved exposure, time to detect local and wider-state mismatch, recovery time, useful task completion and control overhead. Check whether actants understand the distinction between advice and commands. Improvement should be demonstrated without attributing to the AOI benefits already supplied by enforcement or workflow controls.

What an initial demonstration should measure

Add admission pass and false-acceptance rates, restriction activation latency, orientation age at decision time and the share of residual obligations correctly reconciled. Compare local-only context with Grid-enriched context under otherwise equivalent controls. Test whether sharing warnings induces correlated overreaction and whether access restrictions prevent one consumer from obtaining another's private information.

SEI proposition

GUDIYA makes ecosystemic awareness locally actionable. GUDIYA-AOI brings relevant Grid assessment to the actant, while corresponding trusted interfaces inform authorized lifecycle tools. Together they support orientation, continuing security and cross-layer stabilization within a feedback relationship whose effects are verified. SEI proposes no operational attachment without a functioning AOI and demonstrated responsiveness. Implementation evidence must establish where this architecture improves operation and where its limits remain.

Sources and conceptual basis

[1] U.S. Air Force. Luke issues first Generation III helmet.

<https://www.luke.af.mil/News/Features/Display/Article/641647/luke-issues-first-generation-III-helmet/>

[2] Federal Aviation Administration. Airman Education Programs, Spatial Disorientation.

https://www.faa.gov/pilots/training/airman_education/topics_of_interest/spatial_disorientation

[3] NVIDIA Technical Blog, 28 September 2026. NVIDIA Open Agent Safety Platform: A Reference for Continuous In-Silicon Agent Monitoring. <https://developer.nvidia.com/blog/nvidia-open-agent-safety-platform-a-reference-for-continuous-in-silicon-agent-monitoring/>

GUDIYA-AOI, the proposed Grid context service and the integration requirements are SEI architectural proposals. The cited aviation and NVIDIA sources describe their respective systems; they do not validate the proposed GUDIYA implementation.

