117-359-D The Cognitive Strike Package

Mission-Configured Execution, Persistent Headless Actants, and Independent Cognitive Protection

November 2026

This chapter presents a compositional architecture for a multi-actant ensemble supported by the Gudiya Grid and other 'service' actants introduced to ensure that the cognition objectives are achieved while stability is ensured. For inspiration, this compositional architecture derives from the construct used by USAF's Strike Package.

A mission is not a static list of agents. It is a persistent governed undertaking executed by a versioned, dynamically composed, protectively covered system of actants—including long-running, headless, dormant, and cross-mission participants.

Prologue — The Escort That Travels With the Mission

The EA-6B Prowler accompanied aviation strike packages not to perform the strike aircraft's primary task but to supply specialized protection in a contested environment. That analogy led GUDIYA from a simple multi-agent team toward a mission-configured Cognitive Strike Package. The subsequent rock-climber analogy sharpened the boundary: the package is the fully equipped climber; the external GUDIYA Grid is the independent belay. Neither can substitute for the other.

This includes a previously underrepresented class of participant: the persistent, headless actant. Such an actant may remain registered after a user interface closes, sleep between events, awaken weeks later, serve several missions, or reconcile obligations long after a Quarterback's interactive instance terminates. Any execution framework that counts only visible running agents would miss part of its own consequential population.

1. The Architectural Unit: Mission, Package, and Participation

The GCID identifies the governed cognitive mission and survives legitimate changes in agents, hosts, sessions, and execution packages. The Mission Package-ID identifies a specific assembled operational system, whose composition is versioned in a Package Manifest. The Actant-ID identifies a persistent participant. A Participation-ID records the actant's authorized relationship to a mission or package; an Execution Instance-ID identifies a particular runtime incarnation. These identifiers answer different questions and must not be collapsed.

Object	Meaning	Continuity
GCID	Governed undertaking and obligation lineage	Persists through authorized handoffs and settlement
Mission Package-ID	Operational package configuration	May be reconfigured and versioned
Actant-ID	Persistent acting entity	May outlive one host, session, or mission
Participation-ID	Mission-specific membership and authority	Begins, changes, suspends, or expires
Execution Instance-ID	Concrete running incarnation	May start, stop, restart, or scale out

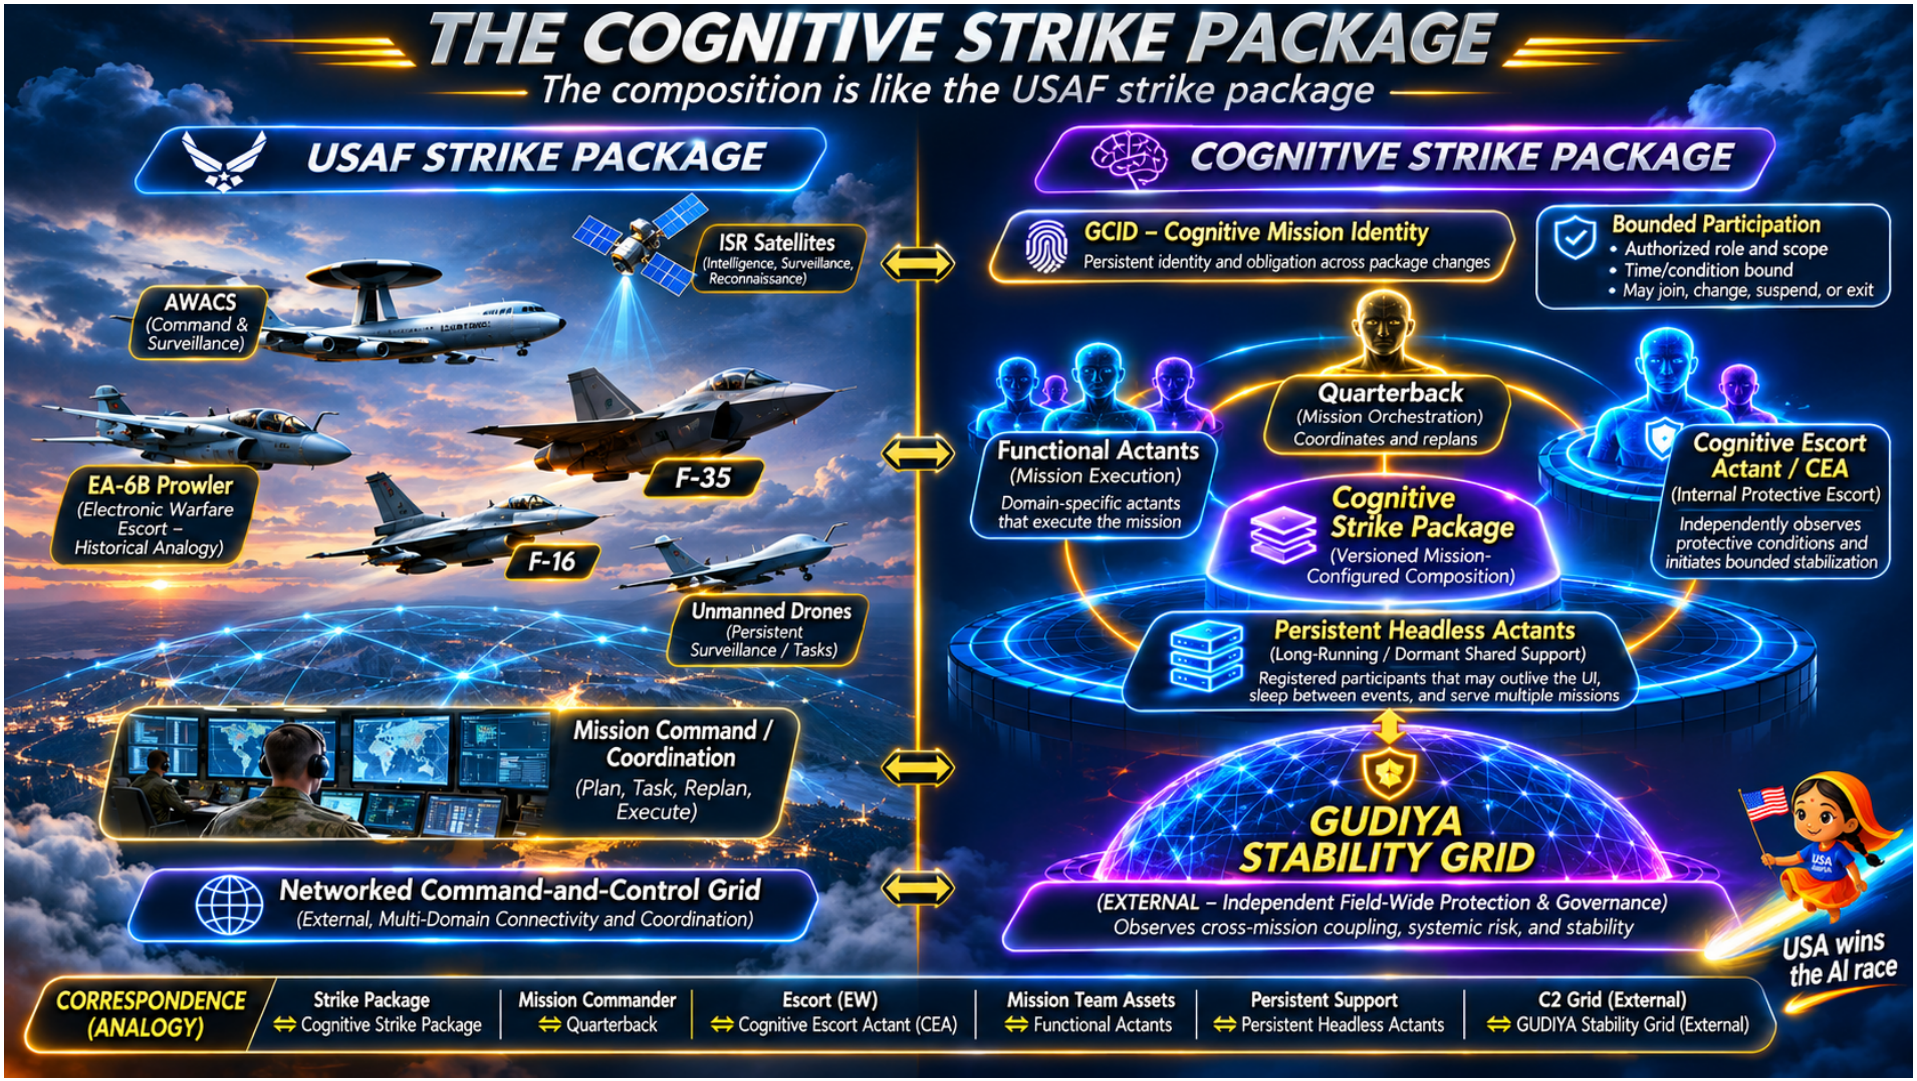
THE COGNITIVE STRIKE PACKAGE
The composition is like the USAF strike package
USAF STRIKE PACKAGE
COGNITIVE STRIKE PACKAGE
ISR Satellites
(Intelligence, Surveillance, Reconnaissance)
AWACS
(Command & Surveillance)
EA-6B Prowler
(Electronic Warfare Escort – Historical Analogy)
F-35
F-16
Unmanned Drones
(Persistent Surveillance / Tasks)
Mission Command / Coordination
(Plan, Task, Replan, Execute)
Networked Command-and-Control Grid
(External, Multi-Domain Connectivity and Coordination)
GCID – Cognitive Mission Identity
Persistent identity and obligation across package changes
Bounded Participation
• Authorized role and scope
• Time/condition bound
• May join, change, suspend, or exit
Quarterback
(Mission Orchestration)
Coordinates and replans
Functional Actants
(Mission Execution)
Domain-specific actants that execute the mission
Cognitive Strike Package
(Versioned Mission-Configured Composition)
Cognitive Escort Actant / CEA
(Internal Protective Escort)
Independently observes protective conditions and initiates bounded stabilization
Persistent Headless Actants
(Long-Running / Dormant Shared Support)
Registered participants that may outlive the UI, sleep between events, and serve multiple missions
GUDIYA STABILITY GRID
(EXTERNAL – Independent Field-Wide Protection & Governance)
Observes cross-mission coupling, systemic risk, and stability
USA wins the AI race
CORRESPONDENCE (ANALOGY)
Strike Package ⇄ Cognitive Strike Package
Mission Commander ⇄ Quarterback
Escort (EW) ⇄ Cognitive Escort Actant (CEA)
Mission Team Assets ⇄ Functional Actants
Persistent Support ⇄ Persistent Headless Actants
C2 Grid (External) ⇄ GUDIYA Stability Grid (External)

2. The Cognitive Mission Triad

The mission has three distinguishable internal functions: the Quarterback coordinates and replans; functional actants execute the mission's work; and the Cognitive Escort Actant (CEA) independently observes protective conditions and initiates only authorized stabilization. These are functions, not necessarily exactly three software processes. The Quarterback must not be able to suppress mandatory protective controls, and the CEA must not silently assume general mission command.

The triad is supported by the external Grid, which observes cross-mission coupling and governs wider-field consequences. A locally coherent package may still destabilize shared enterprise services or physical resources. Internal agreement therefore cannot substitute for field-wide context.

3. The Basic Runtime Population Formula

For a simple package with x initial functional actants, one Quarterback, one CEA, s(t) additional active functional actants spawned during runtime, and r(t) initial functional actants retired, the active actant count is $N(t) = x + 2 + s(t) - r(t)$. This formula is a useful introductory model, but not a complete description of a persistent mission package.

In the general case, the number of active actants is $F(t) + Q(t) + E(t)$, provided each actant is counted once under its assigned role. A larger package may have multiple orchestration and escort actants. An actant may also have several execution instances, and a shared actant may serve several packages; therefore counts of distinct actants, package memberships, and running instances must be reported separately.

4. The Headless Actant Is Not an Exception

A headless actant is a participant without a continuously attended human-facing interface. It may be scheduled, event-driven, persistent, or intermittently executing. Headless does not mean autonomous without authority, and dormant does not mean absent from the mission's responsibility structure. An obligation-monitoring actant can be dormant for hours and then wake to initiate a consequential API call. Its identity and bounded mandate must survive the intervening silence.

The correct abstraction is a persistent actant with execution modes and lifecycle states. A browser session is optional; mission participation and consequential authority are not. A headless actant may serve the original GCID until settlement, or participate in a successor package under an authorized transfer. A cross-mission actant must carry distinct participation and authority bindings for each GCID.

5. Two Population Measures, Not One

Registered population counts the distinct Actant-IDs with valid package participation at a given time, including dormant and nonexecuting headless participants. Executing population counts currently running Execution Instance-IDs. A third useful metric counts active consequential authority grants, which may remain valid even while no instance is running. None of these numbers is interchangeable.

Metric	Definition	Why it matters
Registered actants	Distinct Actant-IDs with valid participation	Protective coverage and mission accountability
Executing instances	Running Execution Instance-IDs	Compute demand and operational load
Authorized consequential participants	Valid, effective authority bindings	Potential to create new commitments

Outstanding obligations	Unsettled commitments linked to the mission	Consequences after execution stops

6. The Versioned Package Manifest

A Package Manifest must record both the planned baseline and authorized runtime changes. The manifest is not merely an agent roster: it is a declarative specification of identities, roles, persistence, authority, protective coverage, and recovery behavior. Every consequential reconfiguration should produce an auditable version or event, while preserving the mission's GCID and historical composition.

Manifest field	Required interpretation
GCID / Package-ID / version	Mission continuity, operational package, configuration lineage
Actant-ID / role / Participation-ID	Persistent entity and bounded mission membership
Execution mode / persistence class	Interactive, headless, scheduled, event-driven; ephemeral or persistent
Execution instance / lifecycle state	Running, dormant, suspended, retiring, terminated
Authority grant / lease	Scope, issuer, expiry, revocation, wake-up conditions
Trigger and resource budgets	Permitted activation events, frequency, concurrency, spending
CEA coverage / Grid attachment	Local protector, escalation route, connectivity state
Checkpoint / recovery policy	Reconstitution, reconciliation, and fitness-to-resume
Outstanding obligations	Open digital or physical commitments and accountable owner

7. Lifecycle States for Persistent Actants

The proposed lifecycle distinguishes Registered, Admitted, Active, Dormant, Suspended, Recovering, Retiring, and Terminated. A dormant actant retains identity but may have no running instance. Suspension removes or narrows authority without erasing its identity or historical obligations. Termination of a runtime instance is not termination of a mission or automatic settlement of its commitments.

State transitions must be triggered by verifiable events, including authorized wake-up, lease renewal, Grid intervention, failed integrity checks, and obligation reconciliation. The model must also handle missed wake-up events, duplicate delivery, host migration, and delayed revocation propagation.

8. Persistent Actant Lease (PAL)

A proposed Persistent Actant Lease (PAL) is a renewable, bounded authority object for long-running or intermittently executing participants. It identifies permitted GCIDs or participation contexts, authorized operations, activation triggers, resource budgets, expiry, renewal rules, and protective coverage. A lease is not an identity and not a substitute for destination-specific credentials or policy evaluation.

A dormant actant may retain its Actant-ID and audit history after its PAL expires, but it should not initiate new consequential actions without valid authority. Wake-up must include an authorization check proportionate to the action and the freshness of the governing context. Lease revocation and local enforcement must be engineered for disconnected conditions.

9. Dynamic Spawning Is a Governed Reconfiguration

A Quarterback may propose spawning additional actants to handle new work. Spawn requests must be bounded by policy, mission budgets, concurrency limits, and CEA coverage capacity. Each new actant or instance must be attributable to a parent delegation and admitted under a distinct identity or explicitly governed replicated identity. Spawned participants must not inherit unlimited authority by default.

The package manifest must preserve the before-and-after topology. A rapidly expanding population can increase throughput but also produce synchronized retries, shared-resource saturation, and cascading obligations. The CEA monitors such changes; the Grid assesses cross-package effects.

10. Escort Coverage Follows Consequential Authority

A critical failure mode arises when interactive work finishes, the Quarterback's session closes, and the escort is deactivated while headless participants continue issuing consequential transactions. The package would appear complete while remaining operationally exposed. Protective coverage must therefore follow active or readily activatable consequential authority, not merely visible execution.

Coverage may be delivered by a dedicated CEA, several subgroup escorts, or a shared protective service with independent enforcement. A dormant actant can be covered by low-overhead monitoring and a guarded wake-up gate. A high-risk actant initiating payment or physical control may require stronger live supervision. Coverage adequacy is mission- and consequence-dependent.

11. Three Scales of Protection

Embedded safeguards protect individual actants through identity checks, bounded tools, policy enforcement, and local constraints. The CEA protects the Intent Group by detecting cross-actant anomalies, unsafe delegation, duplicate commitments, and emerging feedback. The external GUDIYA Grid supplies wider-field observability and authorized cross-mission stabilization. Each scale can fail independently; none should be declared infallible.

12. The Rock Climber and the Belay

The Cognitive Strike Package is the climber with a full complement of competence, coordination, equipment, and local hazard awareness. The Quarterback coordinates the ascent, functional actants perform it, and the CEA supplies embedded protective vigilance. The Grid is the belay: independent stabilization when the package's intrinsic resources are insufficient. A rope does not teach climbing; a capable climber does not eliminate the value of an effective belay.

This analogy also cautions against superficial attachment. A compromised anchor or distracted belayer is not reliable protection. Grid Attachment Integrity must consider freshness of information, intervention readiness, verified authority, and actual connectivity—not a binary connected flag.

13. Disconnected and Degraded Operations

When Grid connectivity is degraded or lost, the package should enter a pre-engineered mode appropriate to mission consequence. Some protective operations must continue, while new irreversible commitments may need to pause. The CEA can enforce local constraints but cannot manufacture missing field-wide context. Reconnection requires reconciliation of intervening actions, leases, obligations, and state before unrestricted authority is restored.

14. Separation of Command and Protection

The Quarterback optimizes mission completion; the CEA monitors protective conditions. Their priorities may diverge. For example, a Quarterback may propose spawning additional purchasing agents to meet a deadline, while the CEA identifies unresolved orders and a rising duplicate-commitment risk. The CEA may request a bounded hold under preauthorized rules, while the Quarterback retains responsibility for replanning. Disagreement is resolved by declared escalation paths, not by whichever model is more persuasive.

15. The Three-Layer AI-Economy

In the revised three-layer model, Layer 1 is the UX Layer with Agentic UX and e-Commerce UX; the paths may exchange control through a Human Commitment Gate. Layer 2 contains enterprise commerce and shared systems of record. Layer 3 contains physical and economic resources and fulfillment. A Strike Package can operate through either UX path, while headless participants may execute without a visible UX session. All consequential actions must still respect enterprise admission and commitment authority.

Cognitive velocity can exceed safe transactional commitment throughput, and physical fulfillment may lag both. The package must not treat a planned purchase, an acknowledged transaction, and a settled physical obligation as equivalent states. GUCL and F-UDL lineage must preserve this distinction across the entire mission lifetime.

16. A Worked Example: The Package That Outlives Checkout

Imagine a procurement mission begins in the AI-Economy.

- It begins with five functional agents, one Quarterback, and one CEA.
- The Quarterback spawns two supplier-comparison agents, then retires one initial agent.
- The interactive phase ends, but a scheduled shipment tracker and a payment-reconciliation actant remain registered under the GCID.
- They are dormant between events and may run only briefly when a shipment status or payment exception arrives.
- The visible team may appear to have finished, yet two consequential authorities and several obligations remain. The package must preserve their leases, wake-up conditions, CEA coverage, and obligation lineage.
- If a supplier acknowledgment is lost and the reconciliation agent attempts a duplicate payment, the local safeguard and escort must be able to block or escalate the action.
- Mission closure is not justified until obligations have an accountable disposition.

17. Small Packages and Federated Packages

For a small Intent Group, a single Quarterback and one CEA may be sufficient. For large populations, the architecture may require multiple Quarterbacks, subgroup escorts, and packages of packages. Shared headless services can serve several GCIDs, but their resource use, authority, and effects must be partitioned or attributed per participation. A single shared monitor must not become an ungoverned cross-mission bridge.

18. Composition Completeness and Cognitive Airworthiness

Preflight qualification must verify that the package possesses mission-appropriate orchestration, execution competence, protective coverage, delegated authority, recovery readiness, and external Grid attachment. Headless and dormant participants belong in the completeness assessment because they can later awaken and act. After significant instability or reconfiguration, Post-Instability Fitness-for-Service Assessment should inspect persistent state, spawned descendants, open obligations, lease validity, and the cause of the original disturbance.

19. GUCL, F-UDL, and Temporal Reconstruction

The cognition ledger and decision/obligation ledger should preserve the time-bounded relationship among GCID, package version, Actant-ID, Participation-ID, Execution Instance-ID, authority grants, and consequences. The Grid must be able to reconstruct who was registered, who was executing, who held authority, which escort covered whom, and what remained unsettled at a specific historical moment. Graph projections may support investigation, but authorization should rely on current authoritative state rather than stale projections.

20. Threats, Failure Modes, and Anti-Patterns

Failure mode	Proposed defense
Orphan headless actant after UI closes	Explicit participation, lease expiry, obligation owner, protective coverage
Spawn storm and uncontrolled population growth	Budgeted delegation, concurrency ceilings, manifest updates
Dormant actant wakes with stale authority	Fresh authorization, policy-version check, re-admission
Shared actant leaks authority across missions	Mission-scoped participation and credential bindings
CEA disappears before obligations settle	Coverage tied to consequential authority and wake-up risk
False Grid attachment confidence	Freshness, heartbeat, intervention path, disconnected-mode controls
Package reconstitution hides contaminated state	Checkpoint integrity, lineage, airworthiness assessment

21. Engineering Demonstrator and Falsification

A controlled demonstrator should compare a conventional orchestrated agent team against a versioned Cognitive Strike Package with a CEA, headless participants, bounded leases, and obligation lineage. Inject lost acknowledgments, wake-up after lease expiry, shared-agent contention, delayed fulfillment, disconnection, and post-intervention restart. Measure unauthorized actions blocked, duplicate commitments, missed obligations, false interventions, recovery time, overhead, and protection coverage gaps. The research hypothesis should be revised if the added architecture does not demonstrate net benefit.

22. Implementation Sequence

Start with a transactional identity and participation kernel, a versioned manifest, and an append-only event history. Add headless lifecycle and PAL enforcement, then CEA coverage rules, cross-layer obligations, and Grid escalation. Keep the initial demonstrator small and compare against established workflow IDs, authorization systems, idempotency controls, and observability tooling. The objective is measurable improvement, not a new name for existing controls.

23. Proposed Doctrines

1. **Mission-Configured Cognitive Execution** : Every consequential Intent Group shall execute through a governed, versioned mission package with explicit identity, composition, authority, protective coverage, and recovery obligations. GCID shall remain distinct from the package configuration.
2. **Persistent Actant Coverage** : A package shall account for interactive, headless, scheduled, event-driven, dormant, and long-running actants. Consequential authority shall remain bounded and protectively covered even when no browser session or execution instance is active.
3. **Separation of Mission Command and Protection** : The Quarterback shall coordinate execution; functional actants shall perform mission work; and the CEA shall independently observe and exercise bounded protective authority. Neither role may unilaterally erase the other's mandated accountability.
4. **The Fully Equipped Cognitive Climber** : The package shall possess the intrinsic capabilities required for its mission, while the external Grid supplies independent stabilization and wider-field governance. Package completeness and Grid Attachment Integrity are complementary conditions of Cognitive Airworthiness.

Epilogue — The Mission Continues After the Screen Goes Dark

- A human may close the browser.
- A Quarterback instance may terminate.
- Functional agents may be retired.
- Yet a shipment may still be in transit,
- A payment may remain disputed, and
- A headless actant may be scheduled to awaken tomorrow.

The Cognitive Strike Package is complete only when the architecture accounts for these continuing identities, authorities, and obligations.

The mission is not the screen. The mission is not the current roster. The mission is the governed undertaking, carried forward through a changing, protectively covered operational system until its consequences are properly resolved.

Closing statement:

- The Quarterback leads.
- The functional actants execute.
- The Cognitive Escort protects.
- The Grid holds the belay.
- And the GCID preserves accountability until consequence closure.