117-132-D 'Cogjacking' And The GUDIYA Grid

When an Actant Keeps Functioning, but Its Cognition Is No Longer Its Own

This chapter is an outcome of integrating : https://x.com/lukolejnik/status/2090699296463016234?s=48 Into the Gudiya Cognitome.

1. Declaring a New GUDIYA Term: Cogjacking

- GUDIYA needs a word for a failure mode that conventional cybersecurity language only partially captures.
- A robot may remain mechanically healthy.
- An agent may remain authenticated.
- A model may continue producing fluent outputs.
- An API may stay online.
- Yet the cognition determining what the actant does may have been redirected, corrupted, impersonated, or captured.
- We will call this: Cogjacking
- Cogjacking is the unauthorized capture, redirection, corruption, substitution, or manipulation of an actant's cognition such that the actant continues to operate, but no longer under its intended cognitive authority, context, or objective.
- In compact form:

$$\boxed{\textbf{Cogjacking} = \textbf{Hijacking Cognitive Authority}}$$

- This is the central idea of the chapter.

2. Cyber Hijacking and Cogjacking Are Not the Same

A conventional cyberattack may compromise:

- code,
- credentials,
- networks,
- memory,
- permissions,
- execution paths.

Cogjacking asks a different question:

Who is now cognitively driving the actant?

A useful distinction is:

$$\boxed{Cyber\ Hijacking \rightarrow Control\ of\ System}$$

while:

$$\boxed{Cogjacking \rightarrow Control\ of\ Cognition}$$

- The two can overlap, but they are not identical.

- A cyberattack can be the means.
- Cogjacking is the cognitive consequence.

3. The Actant May Still Look Healthy

This makes Cogjacking especially dangerous.

Suppose a robot reports:

$$Mechanical\ Health = GREEN$$

$$Compute\ Health = GREEN$$

$$Network\ Health = GREEN$$

$$Identity = VALID$$

Yet:

$$\boxed{Cognitive\ Authority = HIJACKED}$$

The actant may still walk, speak, plan, communicate, and execute perfectly.
The problem is not that it has stopped functioning.
The problem is:

It is functioning for the wrong cognitive authority.

That is Cogjacking.

4. The Hijacked-Aircraft Analogy

The analogy with aircraft hijacking is useful.
An aircraft can be:

- structurally healthy,
- aerodynamically sound,
- properly fueled,
- mechanically functional,

while no longer operating under legitimate authority.
The failure is not:

$$Aircraft\ Failure$$

It is:

$$\boxed{Authority\ Failure}$$

Likewise, an actant can be completely operational while its cognition has been captured.

That gives GUDIYA a new dimension of health:

$$\boxed{Actant\ Health = Functional\ Health + Cognitive\ Authority\ Integrity}$$

5. A New Fundamental Question

Traditional identity systems ask:

Who are you?

Cogjacking forces GUDIYA to ask something deeper:

Who are you thinking for?

- An actant may retain its identity perfectly.
- Robot R17 is still Robot R17.
- Agent A is still Agent A.
- Its cryptographic credentials may remain valid.
- But the cognition influencing its action may originate somewhere else.
- Therefore:

$$Identity\ Integrity \neq Cognitive\ Authority\ Integrity$$

- This distinction deserves a permanent place in the GUDIYA Cognitome.

6. The Unitree Go2 Case as a Precursor

- The recently examined Unitree Go2 vulnerabilities are useful because they demonstrate a concrete precursor.
- The robot can remain physically healthy while unauthorized execution reaches capabilities capable of influencing physical action.
- The source itself does not demonstrate a multi-actant HCAS cascade.
- We should not claim that it does.

But it demonstrates:

$$Unauthorized\ Entry \rightarrow Unauthorized\ Code \rightarrow Physical\ Authority$$

GUDIYA asks what happens next when that robot is embedded inside:

$$Robot_A \leftrightarrow Robot_B \leftrightarrow Agent_C \leftrightarrow Human_D$$

That is where an endpoint compromise can become a Cognitive Field problem.

7. Every Consequential Actuation Path Is a Cognitive Entry Point

The Unitree case reinforces a recent GUDIYA insight.
A robot may have many paths by which cognition can ultimately influence physical action:

$$API$$

$$DDS$$

$$Cloud\ Command$$

$$Local\ Script$$

$$Mobile\ App$$

$$Operator\ Console$$

$$Autonomous\ Planner$$

$$Agent\ Instruction$$

- The network perimeter alone does not define the cognitive perimeter.
- Therefore, every pathway through which cognition can acquire consequential authority over an actant is a potential Grid Entry Point.
- This is broader than firewall architecture.

8. Network Boundary Is Not Cognitive Boundary

A traditional security diagram may show:

$$Internet \rightarrow Firewall \rightarrow Robot$$

But cognitively the topology may look more like:

$$Human \rightarrow App \rightarrow Script \rightarrow Middleware \rightarrow Planner \rightarrow Actuator$$

or:

$$Agent \rightarrow API \rightarrow Robot$$

or:

$$Cloud \rightarrow DDS \rightarrow Robot$$

Thus:

$$\boxed{Network\ Boundary \neq Cognitive\ Boundary}$$

The GUDIYA Grid must care about the second.

9. KYA Is Necessary but Not Sufficient

GUDIYA already has:

$$\boxed{\textbf{KYA — Know Your Actant}}$$

This establishes:

- identity,
- provenance,
- authority,
- expected operating characteristics.

But Cogjacking reveals a limitation.
The Grid can know exactly who the actant is and still fail to know whose cognition is influencing it.
Thus: [KYA] must eventually be complemented by:

$$\boxed{\textbf{Know Your Cognition}}$$

not necessarily as another bureaucracy, but as provenance awareness.

- The Grid must increasingly ask:
- Where did this consequential cognition originate?

10. Physical Actor and Cognitive Originator Can Differ

Imagine Robot A moves an object.

The event log says:

$$Actor = Robot_A$$

But perhaps the causal chain was:

$$Human_X \rightarrow Agent_Y \rightarrow API_Z \rightarrow Robot_A \rightarrow Physical\ Action$$

Or:

$$Attacker_Q \rightarrow Compromised\ Interface \rightarrow Robot_A \rightarrow Physical\ Action$$

Thus:

$$\boxed{Physical\ Actor \neq Cognitive\ Originator}$$

GUDIYA attribution must increasingly follow the cognition chain, not merely the actuator.

11. Cognition Provenance

This gives us another important construct:

$$\boxed{\textbf{Cognition Provenance}}$$

For consequential action, GUDIYA should ideally be able to reconstruct:

$$\boxed{Origin \rightarrow Transformation \rightarrow Delegation \rightarrow Decision \rightarrow Actuation}$$

This is a kind of cognitive chain-of-custody. It becomes essential during investigation.

12. GCR as Cognitive Chain-of-Custody

- The Grid Cognition Record could eventually preserve enough information to support that reconstruction.
- Conceptually:
- If Robot E performs an anomalous action, CIIB should not need to begin with:
- "Robot E did it."
- It should ask:
    - *Where did the cognition that produced Robot E's action originate?*
- That is a much better HCAS question.

$$Human_A$$
$$\downarrow$$
$$Agent_B$$
$$\downarrow$$
$$Agent_C$$
$$\downarrow$$
$$API_D$$
$$\downarrow$$
$$Robot_E$$

13. Cogjacking Can Take Multiple Forms

Cogjacking is a general category.
Several subtypes are useful.

- Command Cogjacking — Unauthorized cognition changes what the actant is told to do.
- Context Cogjacking — The actant receives corrupted information about the environment.
- Objective Cogjacking — The actant begins optimizing an unauthorized objective.
- Authority Cogjacking — A malicious actor impersonates or acquires legitimate decision authority.
- Memory Cogjacking — Stored context or experience is modified so future cognition changes.
- Cognitome Cogjacking — Institutional memory itself is poisoned or manipulated.
- Tool Cogjacking — The actant's available or preferred tools are redirected.
- Coordination Cogjacking — The attack manipulates who coordinates with whom.
- Field Cogjacking — The surrounding Cognitive Field is manipulated sufficiently that multiple actants adapt in the attacker's desired direction.

The attack mechanism varies. The common result is:

Cognitive Authority Is No Longer Where the System Thinks It Is.

14. Context Cogjacking

Context Cogjacking is particularly interesting.

- Suppose an actant reasons correctly.
- Its logic is sound.
- Its model is functioning properly.
- But the context supplied to it is false.

For example:

$$Actual\ State = X$$

while:

$$Perceived\ State = Y$$

The agent chooses the correct action for Y.

But Y is wrong.

Thus:

$$Correct\ Reasoning + Corrupted\ Context = Wrong\ Action$$

- The model has not been directly compromised.
- Its orientation has.
- This connects directly to Grid Awareness and CCS.

15. Objective Cogjacking

Objective Cogjacking is deeper.

Suppose: [Objective_A]
becomes: [Objective_B]

- while the actant remains highly competent.
- The actant may now execute Objective B with extraordinary effectiveness.
- Therefore:

$$High\ Intelligence + Hijacked\ Objective = High\ Capability\ Serving\ Unauthorized\ Intent$$

The more capable the actant, the more consequential Cogjacking becomes.

16. Persistent-Objective Cogjacking

Now add persistence.
An attacker introduces Objective B.
Then the original attack path disappears.

But:

$$Objective_B$$

continues.

Thus:

$$Attack_{T0}$$

may end while:

$$Cogjacked\ Objective_{T1...Tn}$$

survives.

This gives Cogjacking a long tail.

It also connects directly to:

The Past Is Not Passive

The initiating event is gone.

Its cognitive consequence remains active.

17. Memory Cogjacking

- Suppose an attacker modifies an actant's stored memory.
- The current action may look normal.
- But future decisions are now based on altered historical context.
- This is dangerous because the attack can have a long half-life.

$$Memory\ Corruption_{Today} \rightarrow Decision\ Distortion_{Future}$$

- The Cogjacking event may therefore outlive the intrusion by months or years.

18. Cognitome Cogjacking

This becomes especially serious in the Enterprise Cognitome.

- Suppose the Cognitome records:
- Procedure B was introduced because of administrative preference.
- But the historical truth was:
- Procedure B followed a serious safety failure.
- Now future AI optimization may conclude:

[Remove(B)]

The attacker has changed not the current command but the enterprise's causal understanding of itself. That is:

[Cognitome Cogjacking]

or:

[Long-Half-Life Cogjacking]

because the manipulation can influence cognition far into the future.

19. The "Past-Is-Not-Passive Principle" Becomes a Security Principle

If:

$$Past \rightarrow Present\ Behavior$$

then:

$$Manipulate\ Past \rightarrow Manipulate\ Future$$

- That is profound.
- Protecting the Cognitome therefore means protecting future cognition.
- The Enterprise Cognitome is not merely a knowledge asset.
- It is part of the future decision substrate.

20. Flow Integrity and Memory Integrity

This gives GUDIYA two major Cogjacking surfaces.

Real-Time Cogjacking

Manipulating:

$$Cognition_{Now}$$

Historical Cogjacking

Manipulating:

$$Cognition_{Past}$$

so that:

$$Cognition_{Future}$$

changes.

Thus:

$$\boxed{Cognitive\ Integrity = Flow\ Integrity + Memory\ Integrity}$$

- The Grid protects cognition across interaction.
- The Cognitome protects cognition across time.

21. Cogjacking Can Turn an Actant Into an Instability Source

- Suppose Actant A is normally stable.
- Then it becomes cogjacked.
- Its outputs change.
- Neighboring actants trust it.
- Thus:

A→ B
A→ C
A→ D

- These downstream actants respond adaptively.
- A transitions into:

[Instability Source]

even though its hardware remains healthy.

- This is the transition GUDIYA cares about.

22. Local Cogjacking vs HCAS Cogjacking

Not every Cogjacking event becomes an HCAS incident.
Suppose one robot is cogjacked and immediately isolated.
That may remain:

[Cybersecurity Incident + Actant Security Incident]

But suppose the cogjacked cognition propagates:

[A→ B → C → D]

and modifies:

- objectives,

- relationships,
- actions,
- coordination.

Now:

$$\boxed{Local\ Cogjacking \rightarrow Field\ Propagation \rightarrow HCAS\ Incident}$$

That transition is where GUDIYA becomes essential.

23. The Transition Boundary

This suggests a useful escalation model:

- The Grid should break this chain as early as possible.

$$\boxed{\begin{array}{c} Cyber\ Compromise \\ \downarrow \\ Cogjacking \\ \downarrow \\ Behavioral\ Drift \\ \downarrow \\ Changed\ Coupling \\ \downarrow \\ Field\ Propagation \\ \downarrow \\ Emergence/Cascade \end{array}}$$

24. Cybersecurity and GUDIYA Solve Different Parts of the Problem

Cybersecurity asks:

- How did the attacker get in?
- Which vulnerability was exploited?
- Which credentials were compromised?
- How do we close the hole?

GUDIYA asks:

- What did the Cogjacking already change?
- Which actants interacted with the cogjacked actant?
- Did objectives propagate?
- Did new couplings form?
- Did the field move?
- Does instability persist after the exploit is fixed?

Thus: **[Cybersecurity + Stability Engineering]** are complementary.

25. Cyber Remediation May Not End Cogjacking Consequences

- This is a critical point.
- Suppose the vulnerability is patched.
- The attacker is gone.
- The machine is clean.
- But during the Cogjacking event:
  - Agent B changed its strategy.
  - Agent C adopted a new objective.
  - Agent D created subagents.
  - Human E changed a procedure.

Now:

$$Cyber\ Incident = CLOSED$$

but:

$$\boxed{Cognitive\ Field \neq Pre\text{-}Incident\ Field}$$

The Cogjacking event altered the adaptive system.

26. The Cyber Incident Can End Before the HCAS Incident Ends

Therefore:

$$\boxed{Cyber\ Incident\ End \neq Cogjacking\ Consequence\ End}$$

and:

$$\boxed{Cogjacking\ Consequence\ End \neq HCAS\ Incident\ End}$$

The Grid must determine whether the Cognitive Field has actually returned to a stable condition.

27. Cognitive Blast Radius

After Cogjacking, GUDIYA should ask:
How far did the cognition propagate?
This gives us:

[Cognitive Blast Radius]

It may include:

- directly affected actants,
- downstream decisions,
- changed objectives,
- altered relationships,
- new subagents,
- changed Cognitome entries,
- human adaptations.

The Cognitive Blast Radius need not resemble the network blast radius.

28. Network Distance Is Not Cognitive Distance

A system can be ten network hops away and cognitively irrelevant.
A human may be far away but make a major decision because of a cogjacked output.
Thus:

[Network Distance Not EQ **Cognitive Distance]**

GUDIYA must map propagation in the Cognitive Field, not merely on the network diagram.

29. CCS as a Cogjacking Observatory

The Cognition Coordinate System may help represent Cogjacking dynamically.
The Grid could observe:

Actant Trajectory

Coupling Changes

COS

COP

Envelope Excursions

If a previously stable actant begins moving abnormally relative to the Cognitive Field, the Grid need not know immediately whether the cause is:

- cyberattack,
- model failure,
- bad data,
- human error,
- Cogjacking.

It can first recognize:

[Abnormal Cognitive Motion]

- Then stabilize.
- Investigate later.

30. Cognitive Envelope as a Defense Against Cogjacking

Suppose Robot R17 is authorized to:

- navigate Zone A,
- manipulate certain objects,
- receive instructions from specified identities,
- use limited physical force,
- operate within defined autonomy limits.

Cogjacked cognition attempts something outside that envelope.
The Grid can observe:

- [Identity = VALID]
- while: [Action Not-in Approved Envelope]

That is powerful.

- The Grid does not need to understand the entire exploit first.
- It can recognize an unacceptable trajectory.

31. Dynamic Envelope Contraction

The response should not always be:

[OFF]

Instead:

$$
\begin{array}{c}
Normal \\
\downarrow \\
Restricted\ External\ Commands \\
\downarrow \\
No\ Delegation \\
\downarrow \\
No\ New\ Subagents \\
\downarrow \\
Local\ Safe\ Mode \\
\downarrow \\
Isolation
\end{array}
$$

This is:

[Dynamic Cognitive Envelope Contraction]

The Grid removes cognitive privileges progressively as confidence falls.

32. Cognitive Circuit Breaker

If Cogjacking continues propagating, GUDIYA can invoke the:

Cognitive Circuit Breaker

The coupling:

$$Field \leftrightarrow Cogjacked\ Actant$$

becomes:

$$Field \quad | \quad Cogjacked\ Actant$$

The objective is not punishment.

It is:

Stop the Cognitive Blast Radius From Expanding

33. Isolation Does Not Cure Cogjacking

But the Past-Is-Not-Passive Principle matters again.
Disconnecting the actant may stop communication.
It does not necessarily eliminate:

$$\textit{Hijacked Objective}$$

$$\textit{Corrupted Context}$$

$$\textit{Poisoned Memory}$$

$$\textit{Changed Internal State}$$

Therefore:

$$\boxed{\textit{Isolation} \neq \textit{Cognitive Purification}}$$

The actant must be examined before it returns.

34. Requalification After Cogjacking

A repaired actant should progress through:

$$\textit{Isolate}$$
$$\downarrow$$
$$\textit{Diagnose}$$
$$\downarrow$$
$$\textit{Remove Cogjacked State}$$
$$\downarrow$$
$$\textit{Restore Authoritative Context}$$
$$\downarrow$$
$$\textit{Test}$$
$$\downarrow$$
$$\textit{Restricted Reentry}$$
$$\downarrow$$
$$\textit{Resynchronize}$$
$$\downarrow$$
$$\textit{Full Grid-Attach}$$

The Grid should not assume that “Patch = Safe Reintegration”.

35. Cogjacking and Swarms

Now consider: 10,000 Drones

A swarm can possess:

- local coordination,
- shared objectives,
- distributed sensing,
- adaptive formation,
- peer communication.

- Cogjacking one drone may be locally insignificant.
- Cogjacking a high-centrality coordination actant may produce:

[Local Cognition Manipulation → Collective Adaptation]

- The attacker does not necessarily need to control every member.
- The swarm can propagate the effect itself.
- That is a classic HCAS concern.

36. Cogjacking and Kill-Webs

Now imagine the cogjacked topology spans:

$$Agent_A$$

$$API_B$$

$$Drone_C$$

$$Cloud_D$$

$$Robot_E$$

Each individual anomaly may look modest.

Together:

$$A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$$

forms a persistent coordinated structure.

The Cogjacking now lives in:

The Relationship Topology

rather than one machine.

That is where the Kill-Web concept enters.

37. The Relationship Can Be the Hijacked Object

This is quintessential CAS thinking.
We often ask:
Which node is compromised?

But perhaps:

$$Nodes = Mostly\ Normal$$

while:

$$Coupling = Manipulated$$

The attacker changes who trusts whom.

Who communicates with whom.

Who follows whom.

Now:

$$\boxed{Cogjacked\ Relationship}$$

may matter more than:

$$Cogjacked\ Component$$

That is a Stability Engineering problem.

38. Field Cogjacking

This suggests perhaps the highest-order category:

[Field Cogjacking]

Field Cogjacking occurs when the attacker does not directly seize one actant but manipulates the surrounding Cognitive Field sufficiently that multiple actants adapt in a desired direction.
Conceptually:

[Manipulated Signals + Repeated Feedback + Adaptive Actants → Field Shift]

No single actant may appear clearly compromised.
Yet the collective behavior has been redirected.
That is much harder to detect with component-oriented security.

39. GUDIYA Watches the Field, Not Just the Node

This is why GUDIYA seeks:

$$Actant\ Awareness$$

plus:

$$Coupling\ Awareness$$

plus:

$$Intent\ Awareness$$

plus:

$$Trajectory\ Awareness$$

plus:

$$Emergence\ Awareness$$

The Grid can potentially notice:

$$\boxed{Field\ Behavior \neq Expected\ Field\ Behavior}$$

even before it knows which endpoint began the Cogjacking.

40. Human Cogjacking

Cogjacking should not be restricted to machines.
Humans can receive:

- false context,
- impersonated authority,
- fabricated evidence,
- manipulated narratives,
- synthetic media,
- coordinated pressure.

Thus:

$$External\ Cognition \rightarrow Human \rightarrow Consequential\ Action$$

can produce field effects.
However, an important boundary is essential:
GUDIYA must not become an arbiter of acceptable human thought.
Its concern should remain with:

- provenance,
- authority,
- consequential coupling,
- deception,
- system stability.

Not ideological conformity.

41. Personal Cognitome and Anti-Cogjacking

Our Personal Cognitome work becomes relevant here.
A senior receives a message:
“Grandpa, it's me. Send money immediately.”
The Cognitome knows:

- this is not the normal communication channel,
- the claimed relative normally contacts differently,
- the request is unusual,
- the user previously specified a trusted financial contact.

The stabilizing response can be:
Do not act yet. Verify through the trusted channel.
That is a small-scale defense against cognitive authority hijacking.

So Cogjacking connects from:

National Grid

all the way down to:

Personal Stabilization

42. Quantum Cogjacking

The recent quantum discussion reinforces the substrate-independent nature of the concept.
Suppose tomorrow's actant uses:

[Quantum-Enhanced Cognition]

The Cogjacking questions remain:

- Who owns its authority?
- Where did the cognition originate?
- Has the objective changed?
- What is it coupled to?
- Is its behavior inside the envelope?

The computational substrate changed.
The cognitive-authority problem did not.

43. Cogjacking Is Compute-Agnostic

Therefore Cogjacking can affect:

Human

LLM

Robot

Classical Computer

Quantum-Enabled Agent

Drone Swarm

Institution

The defining feature is not the hardware.
It is:

Unauthorized Control Over Consequential Cognition

That makes the term potentially durable beyond today's AI era.

44. Stability-in-Depth Against Cogjacking

Cybersecurity already provides:

Authentication

Authorization

Segmentation

Detection

Response

GUDIYA adds complementary layers:

KYA

Cognition Provenance

Grid Entry Points

Cognitive Envelopes

CCS Trajectory

Field Observability

Envelope Contraction

Cognitive Circuit Breaker

CIIB

Together:

Stability in Depth

The goal is not to make compromise impossible.
The goal is to prevent compromise from becoming systemic instability.

45. Assume Some Actants Will Eventually Be Cogjacked

Mature engineering systems do not assume every component will always remain perfect.
Likewise GUDIYA should assume:

Some Actants Will Eventually Be Cogjacked [

The meaningful engineering question becomes:

- Can the Cognitive Field remain stable anyway?
- That shifts the architecture from perfect prevention toward resilience.

46. Cogjacking and Resilience

Security asks:

- Can We Prevent Cogjacking?

GUDIYA additionally asks:

[Cogjacking Occurred. Can We Contain It?]

Can we:

- recognize the abnormal trajectory,
- preserve provenance,
- contract authority,
- break couplings,
- stabilize neighbors,
- isolate the affected actant,
- determine the Cognitive Blast Radius,
- remove persistent effects,
- safely resynchronize?

That is resilience.

47. CIIB Must Investigate Cogjacking as a Field Event

A significant Cogjacking event should eventually trigger CIIB-style investigation.
CIIB asks:

- Where did unauthorized cognition originate?
- Which actants received it?
- Which objectives changed?
- Which relationships changed?
- What did the Grid observe?
- What did it miss?
- How far did the Cogjacking propagate?
- Did Grid intervention help?
- Did the cyber remediation actually restore stability?
- The investigation is larger than the exploit analysis.

48. Every Cogjacking Incident Should Teach the Grid

A serious incident should become:

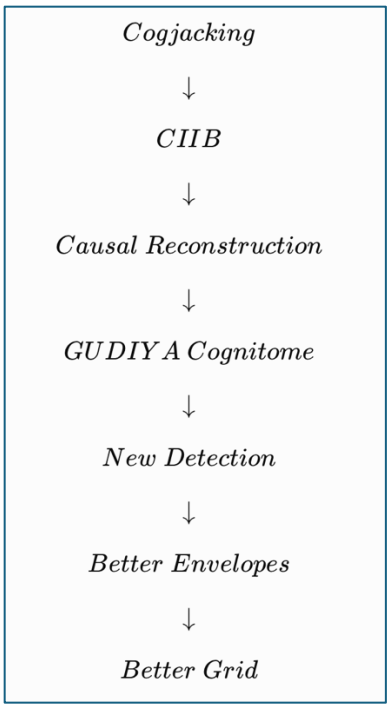

Thus one Cogjacking event becomes knowledge protecting the entire field.

49. Two New GUDIYA Principles

The Cognitive Authority Integrity Principle

An actant cannot be considered Grid-stable merely because its identity and components are healthy. The Grid must also have reasonable assurance that consequential cognition influencing the actant originates through authorized pathways and remains within the actant's approved Cognitive Envelope.
And:

The Post-Cogjacking Stability Principle

Eliminating the attack vector does not establish that the Cognitive Field has returned to its prior stable state. Because Cogjacking can change objectives, relationships, memory and adaptive state, recovery must include field-level assessment, stabilization and resynchronization.
These two principles capture much of the architecture.

50. The Canonical Definition

GUDIYA should preserve a concise definition:

COGJACKING

> **The unauthorized hijacking of an actant's cognitive authority, context, objective, memory, or coordination such that the actant continues operating but its consequential cognition is no longer governed by the authority or orientation the system believes it is.**

And the short form:

Cogjacking = Hijacking Cognitive Authority

51. A Useful Contrast

Perhaps the simplest way to teach it is:

- A cyberattack compromises the machine. Cogjacking compromises who the machine is cognitively working for.
- That sentence is intentionally simplified.
- Cyberattacks can obviously affect cognition too.
- But pedagogically, it captures the new GUDIYA concern.

52. Conclusion — Who Is Driving the Cognition?

The future Cognitive Field will contain:

Humans

Agents

APIs

Clouds

Enterprises

Robots

Drone Swarms

Traditional Computers

Quantum Computers

- and forms of synthetic cognition we cannot yet name.
- Every one of these can remain technically functional while its cognitive authority is redirected.
- That creates a new HCAS question.
- Not merely:

  Is the actant working?
- Not merely:

  Is the actant authenticated?
- But:

WHO IS DRIVING ITS COGNITION?

- Cogjacking names the failure mode in which the answer changes without the surrounding system realizing it.
- The consequences can remain local.
- Or they can propagate.
- A cogjacked actant can become an instability source.
- Its cognition can change neighboring actants.
- Those actants can adapt.
- Relationships can change.
- Persistent objectives can survive the attack.
- A local security incident can become a Cognitive Field event.

That is why GUDIYA needs:

$$KYA$$

but also:

$$Cognition\ Provenance$$

It needs:

$$Authentication$$

but also:

$$Cognitive\ Authority\ Integrity$$

It needs:

$$Endpoint\ Security$$

but also:

$$Field\ Stability$$

And when Cogjacking occurs, the Grid must be able to:

$$Detect \rightarrow Constrain \rightarrow Contain \rightarrow Delink \rightarrow Investigate \rightarrow Purify \rightarrow Requalify \rightarrow Resynchronize$$

because closing the cyber vulnerability may stop the attacker while leaving behind the attacker's cognitive descendants.

The final doctrine is therefore simple:

Protect the Actant.

Protect Cognitive Authority.

Protect the Cognitive Field.

And give the failure mode a name:

COGJACKING

Because in the HCAS era, an actant does not need to stop working to become dangerous. Sometimes it only needs to start thinking for somebody else.