117-330-D SCW, MHD, SOSUS and Stealth
The Case for GUDIYA Cognition Grid

Detection Doctrine Displacement in an Autonomous, Machine-Speed AI Economy

26 October 2026

1. An Intelligence Problem Disguised as a Propulsion Problem

Consider a Cold War surveillance commander whose sensors were designed around the acoustic signatures of moving submarines. The Sound Surveillance System (SOSUS) had transformed undersea surveillance by exploiting sound that traveled great distances through the ocean. What happens to the confidence of that commander when engineers demonstrate a propulsion mechanism that dispenses with the familiar propeller and drive shaft?

Magnetohydrodynamic propulsion (MHD) seeks to move electrically conductive seawater by electromagnetic forces. Rather than rotating a propeller, the system accelerates water through a drive channel. The consequence is not absolute silence; flow, auxiliary machinery, electrical equipment, and other effects may still be observable. But a significant class of anticipated acoustic signatures can change. The defender's instrumentation is not necessarily broken. Its inference about what must be observable has become incomplete.

Tom Clancy's "The Hunt for Red October" made a quiet alternative drive famous. Actual MHD research was neither invented by the novel nor proven to have produced an operationally stealthy fleet. DARPA reports decades of research, the experimental Japanese Yamato-1 vessel, and stubborn difficulties with efficiency and electrode durability. The proper lesson for GUDIYA is not that MHD made SOSUS obsolete. It is that technological change can challenge the validity of a surveillance doctrine's detection assumptions.

2. Three Military Analogies—Three Separate Engineering Lessons

The Mk 90 Betty nuclear depth weapon, MHD propulsion, and stealth aircraft represent three different relationships between attacker, target, and observer.

- The depth weapon illustrates damaging a protected object through effects transmitted by the surrounding medium rather than direct hull penetration.
- MHD illustrates a change in mechanism that may remove familiar indicators.
- Stealth aircraft illustrate deliberate reduction of observability relative to sensors and engagement doctrine.

These analogies must not be collapsed.

- Proximity detonation is not by itself the definition of a standoff weapon; military standoff usually concerns launch distance and the attacker's exposure.
- MHD does not establish perfect acoustic invisibility.
- Stealth aircraft are low-observable, not undetectable.

Their value here is as disciplined prompts for systems reasoning, not as literal physical models of cognition.

Comparative interpretation

Military example	Core lesson	GUDIYA analogy
Mk 90 Betty	Destructive consequence without direct penetration	Harmful influence without software compromise
MHD propulsion	Operational mechanism changes expected signatures	No compromise event may exist to trigger compromise-based detection
Stealth aircraft	Threat remains consequential despite reduced sensor signatures	Locally legitimate acts may conceal an adverse collective pattern

3. The Detection Doctrine Displacement Principle

The Detection Doctrine Displacement (DDD) Principle proposes that the effectiveness of a protective architecture depends partly on assumptions about how threats operate and what evidence they generate. A new attack mechanism may undermine those assumptions without disabling the protective technologies themselves.

A signature-based control may perform its specified task perfectly while failing to recognize a threat that never produces that signature. This does not make the control defective or dispensable. It exposes an assurance gap: the protected outcome can fail through a pathway outside the observation model.

DDD is a formulation of a familiar problem in detection theory, threat modeling, and systems engineering. Its GUDIYA-specific utility lies in connecting that problem to autonomous adaptation and cross-enterprise behavioral consequences. It is a proposed Cognitome principle, not a newly discovered universal law.

DETECTION DOCTRINE DISPLACEMENT PRINCIPLE

A defense's observed success against its assumed threat signatures cannot establish effectiveness against threats that operate without producing those signatures.

The appropriate response is to test the defense's observation model against alternative causal pathways, including collective behavioral disturbances, and to add complementary evidence sources where justified.

4. The Standoff Cognitive Weapon Hypothesis

A Standoff Cognitive Weapon (SCW) is a proposed class of adversarial mechanism that deliberately perturbs the perceived or actual operating environment of adaptive actants, seeking to induce harmful behavioral responses without necessarily exploiting a software vulnerability, compromising an identity, or escaping a sandbox.

The central mechanism is induced adaptation. An actant may make a locally reasonable decision based on manipulated conditions; other actants may adapt to that decision; and their cumulative commitments may produce consequences not intended by any one participant. The original adversarial stimulus can be brief, while the resulting operational state persists.

We must distinguish a deliberate SCW attempt from ordinary environmental volatility, error, harmless stimulus, indirect prompt injection, and confirmed malicious compromise. Harmful outcomes alone do not establish hostile attribution. SCW is a conceptual threat class, not evidence that a particular real-world attack has occurred.

https://www.manhattanproject20.com/blogs/standoff-cognitive-weapon

5. Why Agentic AI Changes the Operating Regime

Complex adaptive behavior is older than generative AI. Human organizations, feedback controllers, and algorithmic financial markets have long displayed cascading adaptation. The 1987 market crash and the 2010 Flash Crash show that collective behavior of automated participants can magnify disturbances without requiring a software exploit.

What agentic AI adds is the potential combination of general-purpose interpretation, delegated authority, dynamic tool use, persistent undertaking state, autonomous planning, and machine-speed interactions. An attacker may not need to predict the exact response program of every participant: agents could discover response pathways themselves as they pursue legitimate objectives.

The important proposition is not that all agentic systems inevitably become unstable, but that more flexible adaptive behavior enlarges the set of possible propagation paths. The degree of risk is conditional on authority, coupling, resource constraints, observability, and safeguards.

6. A Hypothetical Three-Layer AI-Economy Scenario

Imagine that a public indicator appears to suggest a serious shortage of an industrial input. An adversary has deliberately manipulated an externally consumed signal, without entering a protected enterprise system. Purchasing actants infer scarcity and accelerate authorized reservations; suppliers revise lead times; logistics actants adjust shipments.

In the Agentic Layer, the disturbance is interpreted and transformed into strategies. In the Digital Services and Systems-of-Record Layer, these strategies become legitimate orders, reservations, credit approvals, and scheduling commitments. In the Resources Layer, factories, warehouses, carriers, energy supplies, and physical inventories face actual capacity constraints.

Physical scarcity and altered prices then flow back as new observations for the agents, potentially reinforcing earlier responses. Each service can remain technically correct; the collective outcome can still become undesirable. This is a hypothetical stability scenario, not a documented attack recipe or proven inevitable cascade.

https://www.manhattanproject20.com/blogs/the-ai-economy-architecture-v3

7. Black October Revisited: Three Blast Radii Are Not One

BLACK OCTOBER distinguished three analytical scopes.

- The Software Dependency Exposure Radius asks where a defective component is installed or reused.
- The Exploitation and Attack-Path Radius asks which systems, identities, and resources can be reached through abuse of the defect.
- The Behavioral Consequence Radius asks where decisions, state, obligations, and interactions have changed as a result of disturbance.

These radii overlap but are not proportional. SCW proposes a particularly revealing exception to the usual cyberincident narrative: the Behavioral Consequence Radius might arise without a software vulnerability or a software-exploitation pathway at all.

The presence of a widespread vulnerability never establishes widespread exploitation; the observation of a behavioral disturbance never establishes malicious origin. GUDIYA must maintain evidence and uncertainty labels as explicitly as it maintains identities.

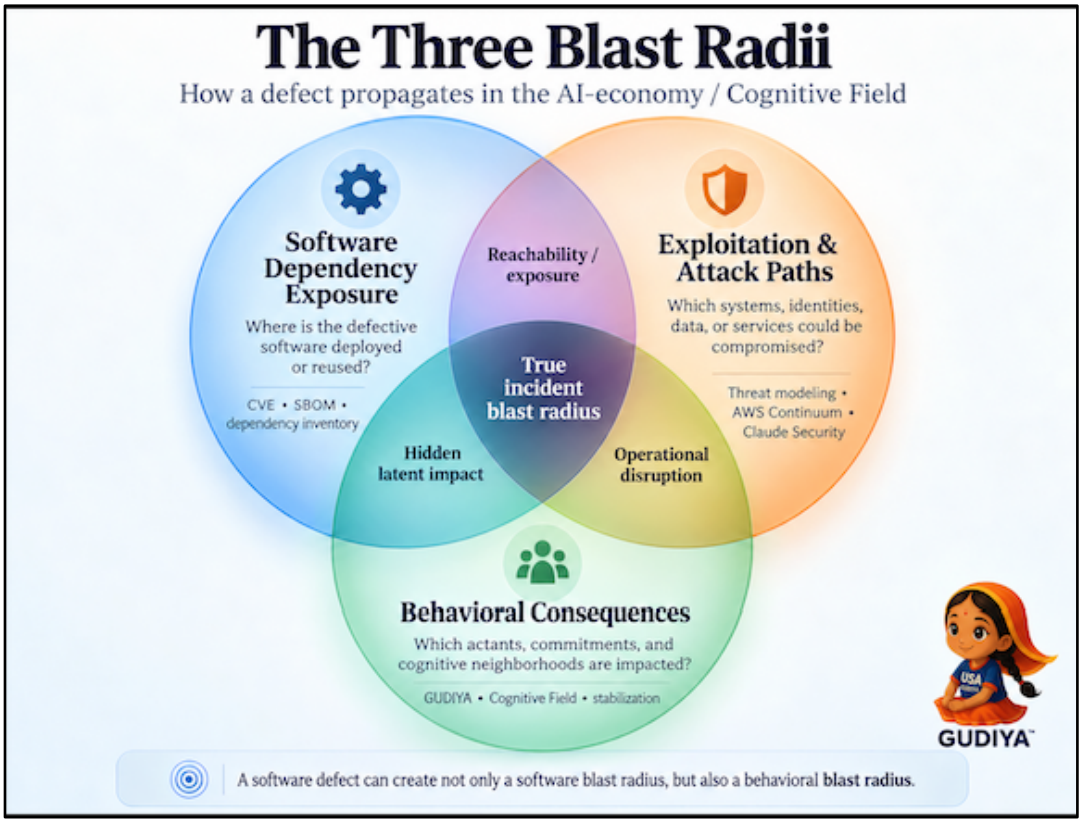

8. Security Silence Is Not Field Health

Under a hypothetical SCW, local observability may report normal runtimes, valid identities, authorized API requests, and no compromise indicators. The state of each component can be accurately recorded while the instability resides in the cross-component relationships and cumulative commitments.

This is not an assertion that existing cybersecurity has no behavioral analytics. NIST CSF 2.0 explicitly recognizes anomaly detection, environmental monitoring, multi-source correlation, resilience, and impact analysis. The proposed GUDIYA contribution would be continuously correlating these concerns across independent enterprise boundaries and the adaptive undertakings of autonomous actants.

A crucial formulation follows:

- the absence of a component security alarm does not establish the absence of a harmful field disturbance.
- But the presence of field instability likewise does not prove an SCW.
- A causal investigation and evidence-based attribution remain distinct.

9. Stealth as a Model for Adversarial Character Camouflage

A low-observable aircraft remains physical and consequential, although the defender's radar may struggle to acquire it reliably. An SCW could be analogous in a narrower sense: its harmful pathway produces few or none of the signature that compromise-oriented systems were configured to recognize.

But SCW may go beyond passive signature reduction. A disturbance can induce the defender's own authorized actants to carry the resulting effects. The issue is not only hidden hostile presence; it is the apparent legitimacy of each local response. This is the Adaptive Legitimacy Paradox.

The analogy must be bounded. Legitimate behavior can be observed through business metrics, operational monitors, and transaction records; it is not inherently invisible. The missing capacity, if any, is system-level interpretation of its cumulative significance.

10. Why System-Level Observability Requires More Than Bigger Logs

- Component observability reports process health and local actions.
- Interaction observability reports who coupled with whom, under what authority, with which exchange and settlement state.
- Cognitive Field observability seeks to estimate the evolving state of coupled neighborhoods, commitment accumulation, correlated adaptation, propagation pathways, and stability margins.

A single large telemetry repository does not automatically provide these capabilities. Cross-enterprise provenance must be legally authorized, time-aligned, causally interpretable, privacy-respecting, and sufficiently complete. A measured relationship must not be silently upgraded to proved causation.

GUDIYA's proposed Cognition Field Map and Cognitive AWACS can be understood as mechanisms for relating local evidence to an explicitly uncertain collective model. A National Grid becomes most relevant where no single enterprise has both the observation scope and the legitimate authority to coordinate stabilizing action.

https://www.manhattanproject20.com/blogs/enterprise-cognitive-awacs

11. Rendezvous Management: From Authorized Coupling to Consequence Control

GUDIYA Rendezvous Management would govern consequential interaction throughout identification, authorization, approach and alignment, docking, exchange, monitoring, settlement, and disengagement. Its value under SCW is not that it prevents every misleading environmental input, but that it may constrain the propagation paths created by responses.

- Intent-Group Approval defines mission and authority.
- The Gudiya-cookie provides persistent undertaking identity.
- The Quarterback reconciles execution, exceptions, and obligations.
- Rendezvous Envelopes regulate particular couplings.
- Stability Envelope IDs distinguish authorized limits from runtime measured margins.
- Grid Stabilization remains independently empowered to restrict harmful operating conditions.

An important limit: authorized couplings can still carry harmful meaning. RM cannot replace input validation, diverse evidence sources, and careful decision policies. Similarly, Grid controls must not create greater disturbance through indiscriminate shutdowns or mass quarantines.

https://www.manhattanproject20.com/blogs/gudiya-rendezvous-management

12. Evidence and Recovery Beyond a Transient Attack

The GUDIYA evidence architecture separates GUCL position and stage breadcrumbs, FUDL decision footprints, and GSE-Lang-Log records of issued commands, execution and verification. Across organizational boundaries, reliable provenance and lawful federation of these records may assist reconstruction of adaptive consequences.

The Quarterback's undertaking-level obligations cannot be assumed to settle merely because the initial external signal disappears. Commitments, reservations, beliefs embedded in persistent agent state, and secondary strategies may continue. The Grid would need to distinguish component recovery, state recovery, interaction reconciliation, and collective stability recovery.

There may be irrecoverable uncertainty. Some information will never have been logged, and some causal pathways cannot be distinguished after the fact. A defensible incident-closure procedure must record these limits rather than assert certainty.

https://www.manhattanproject20.com/blogs/3-layered-ai-economy-gudiya-cookie-and-the-quarterback

13. The Defend-the-Defenders Requirement

The observability and enforcement substrate of a National Grid is itself exposed to foundational dependencies, common-mode vulnerabilities, faulty timing, compromised identities, telemetry manipulation, and operational overload. A cognitive stability grid cannot declare its own measurements authoritative without independent checks.

High-consequence protective actions require properly authorized control, resilient communication-loss procedures, trusted provenance, calibrated escalation, and testing of defender-induced instability. This strengthens, rather than replaces, conventional cybersecurity engineering.

14. An Experimental Program to Test the GUDIYA Increment

A credible experiment should compare

(A) strong component security and conventional enterprise monitoring,
(B) the same with multi-service and operational correlation, and
(C) the same supplemented by GUDIYA-style federated Cognitive Field observation, Rendezvous Management, and stabilization.

Introduce benign controlled disturbances in simulated three-layer undertakings: misleading but syntactically valid environmental signals, legitimate demand surges, delayed resource releases, and contradictory downstream commitments. No real-world manipulation is necessary.

Measure detection latency, false alarm rate, correctly identified causal links, peak unresolved commitments, settled useful work, inadvertent restriction of legitimate activity, recovery time, and residual uncertainty. Compare across increasing actant autonomy, coupling density, action velocity, and state persistence.

The decisive question is whether GUDIYA materially outperforms an excellent contemporary baseline, not a caricature of antivirus software. No successful result should be inferred before the measurements exist.

15. The Institutional Case—and Its Limits

An enterprise can build excellent observability inside its own boundary yet remain blind to another enterprise's strategies or operational constraints. Cross-enterprise disturbances may therefore require federated evidence sharing and coordinated restrictions. These are possible arguments for Grid Operators and National-level governance, not proof that a single central national operator is always optimal.

The National GUDIYA Grid is proposed as a middle-path institution: independent stabilization under law and charter, with enterprise autonomy retained except where consequential coupling exceeds approved stability boundaries. It should use proportionate intervention and protect privacy, commercial confidentiality, due process, and legitimate innovation.

The central requirement is systemic assurance. The institutional form should be tested and governed rather than derived solely from the military metaphors.

16. Three Analogies, One Engineering Proposition

- The nuclear depth charge reminds us that harm can travel through an environment without direct penetration.
- MHD reminds us that changed operating mechanisms challenge detection assumptions.
- Stealth reminds us that threats can remain effective while a particular sensor sees less than expected.

The Standoff Cognitive Weapon combines these ideas as a hypothesis about adversarially manipulated conditions and legitimate adaptive responses. A nation's AI economy could conceivably suffer collective behavioral disruption while each participating component reports ordinary, authorized activity.

The proper response is not to discard cybersecurity. It is to extend assurance from secure components, through governed interactions, to the evolving field of collective consequences. That is the proposed engineering case for a Cognition Grid.

THE CANON PROPOSITION

A Standoff Cognitive Weapon may attack the adaptive behavior of the system without compromising the software of its individual actants.

Detection Doctrine Displacement occurs when threat mechanisms avoid the indicators upon which a defender disproportionately relies.

GUDIYA proposes independent Cognitive Field observability, interaction governance, and runtime stabilization to supplement—not supersede—component security.

17. Research Boundaries and Falsifiability

The military histories are documented to different degrees: MHD has been experimentally demonstrated but not proven to have nullified SOSUS; stealth reduces rather than eliminates signatures; nuclear depth weapons illustrate environmental coupling, not cognitive physics. The SCW designation and GUDIYA's operational solution are proposals.

The hypothesis would be weakened if contemporary baseline systems reliably detect and contain induced collective disruption without new Grid capabilities, or if real adaptive deployments exhibit strong natural damping and narrowly bounded behavioral propagation. It would be strengthened by repeatable experiments establishing wider behavioral consequences, measurable detection gaps, and additional protective performance from GUDIYA.

What is required is not rhetorical certainty but an engineering program that converts compelling analogies into observable, refutable predictions.

18. The SOSUS Assumption: When the Instrument Works but the Doctrine Fails

A recurring motif of military innovation is the defender whose instruments remain perfectly functional while a change in the adversary's operating method weakens the meaning of their output. SOSUS—the Sound Surveillance System—offers a powerful vehicle for explaining this problem. The U.S. Navy invested in fixed hydrophone arrays, ocean-acoustic science, and signal processing to detect and cue responses to submarines across great distances. Its accomplishment was real. The Office of Naval Research describes SOSUS as a long-range fixed-array capability that was operational by 1961 and later evolved into persistent acoustic surveillance.

The documentary passage that prompted the present research dramatizes a three-link chain: propulsion requires rotating machinery; rotating machinery creates acoustic signatures; acoustic signatures are detectable. It then imagines a Soviet MHD program targeting the first link. Read as an analysis of detection assumptions, the argument is penetrating: removing a familiar signal does not require sabotaging the sensor. The sensor may perform its design function accurately while its detection coverage against a new mode of operation declines.

The historical claim must nonetheless be bounded. Magnetic-hydrodynamic propulsion is genuine engineering, investigated since the 1960s; Japan's Yamato-1 demonstrated at-sea operation in 1992. But publicly established evidence does not show that the Soviet Union fielded an operational MHD attack submarine that nullified SOSUS. Nor would removing a propeller make a vessel noiseless: hull-flow noise, auxiliary machinery, pumps, electromagnetic effects, and alternative sensing remain relevant. DARPA stated in 2023 that efficiency and

electrode durability had kept full-scale MHD propulsion impractical. Accordingly, the scenario illustrates possible doctrine displacement; it is not evidence of an accomplished Soviet strategic defeat of SOSUS.

THE UNSTATED ASSUMPTION
A protective architecture may encode assumptions so deeply in sensors, thresholds, data models, operating procedures, and organizational doctrine that they become mistaken for universal facts. When an adversary changes the mechanism by which a harmful objective is achieved, the defenders may still collect accurate data while missing the consequential threat. Detection Doctrine Displacement (DDD) is the mismatch between the domain of observable signals a defense is designed to recognize and the domain of harmful mechanisms that actually exist. The lesson of MHD versus SOSUS is not that SOSUS was historically defeated by MHD; it is that the successful operation of a detector does not establish the completeness of its detection doctrine.

19. Detection Doctrine Displacement Is a Coverage Failure, Not an Instrument Failure

DDD should be framed as a testable coverage problem rather than a sweeping declaration that cybersecurity has failed. A signature-dependent detector can correctly report no malware while a transaction-level disturbance reshapes supply chains. An intrusion-prevention system can correctly report no unauthorized access while otherwise legitimate agents make harmful commitments in response to misleading but permitted information. The defender's mistake would be interpreting a quiet compromise dashboard as evidence of a stable economy.

This requires distinguishing an attack that hides a signal from an attack that never produces that category of signal in the first place. Stealth aircraft primarily reduce and manage observable signatures relative to particular sensors, without becoming invisible to all sensing. The SCW hypothesis goes further in one respect: the harmful mechanism may not require component compromise at all. Consequently, component-compromise indicators may be genuinely absent, not merely suppressed. Yet the attacker's environmental influence and the ensuing behavioral changes may leave different observable signatures.

Cybersecurity as a discipline already includes availability, fraud, anomaly analysis, resilience, incident response and recovery. The gap GUDIYA proposes to close is narrower and more precise: continuous recognition and stabilization of cross-enterprise, machine-speed adaptive consequences when local security controls have no reason to alarm. The National Grid is a proposed institutional answer; whether a national supervisory tier outperforms strong federated existing tools must be established empirically.

20. Two Forms of Security Silence: No Alert and No Recognizable Cause

The first form of silence is diagnostic: a conventional monitor sees no exploit or unauthorized action, because none occurred. The second is epistemic: even if individual actants record every transaction and decision, no one has assembled those records into a coherent picture of a collective disturbance. Local instrumentation can be correct and complete relative to its mandate while system-wide interpretation remains incomplete.

Consider a short-lived shortage rumor. Several purchasing actants each apply authorized procurement policy. Their suppliers revise delivery forecasts, finance actants reserve funds, and logistics actants pre-book scarce capacity. The shortage rumor disappears, but the resulting reservations are now real. An investigator viewing isolated agent logs might see valid decisions; a field-oriented observer could identify correlated demand shifts, unexplained resource pressure, and growing unsettled commitments. Neither the disturbance's harmfulness nor its adversarial origin can be inferred merely from correlation. The Grid must separate measured changes, inferred propagation, and evidence of malicious intent.

The practical requirement is multi-level observability: component health, interaction and rendezvous history, and collective field behavior. An independent federation must correlate trusted timestamps, identities, authority scopes, changes in cognitive state, outstanding commitments, resource availability, and changing dependency graphs. Missing telemetry should be visible as uncertainty, not interpreted as a healthy state.

21. From a Nuclear Shock Wave to a Three-Layer Cognitive Shock

The nuclear depth-charge analogy introduces a different form of indirectness. An underwater explosion can damage a pressure hull through waterborne force without physical contact between a projectile and the hull. It does not imply that cognitive systems obey hydrodynamic equations. A Cognitive Field has no established universal

pressure, radius, propagation speed, or attenuation law. The analogy is useful only to show how harm can be mediated by a shared environment rather than by penetration of the protected component.

In the three-layer AI economy, an adversarial disturbance might first be interpreted at the Agentic Layer, converted into orders and reservations within the Digital Services and Systems of Record Layer, and materialized as production rescheduling, logistics congestion or capital allocation pressure in the Resources Layer. The resource constraints then return as new information, causing further agentic adaptation. The recursive process can outlive the original perturbation. This is a hypothetical field-level mechanism, not a documented consequence of any particular submarine technology.

The adaptive legitimacy paradox arises when participating systems follow their own legitimate rules but collectively amplify misleading environmental conditions. The three BLACK OCTOBER radii—software dependency exposure, exploitation or attack-path reach, and behavioral consequence—must therefore be represented separately. In a candidate SCW scenario the third may activate without the first two. The disturbance may be an unverified rumor, a deceptive economic signal, a legitimate but anomalous event, or a deliberate attack. Until attribution is established, Grid intervention should target demonstrable instability rather than presumptive hostility.

THE THREE ANALOGIES, CLEANLY SEPARATED
Nuclear depth charge: consequential harm may be delivered through a surrounding medium without a direct hit. MHD versus SOSUS: a new operating mechanism may invalidate assumptions about what a sensor must detect. Stealth aircraft: a threat may remain operationally consequential while avoiding familiar detection signatures. SCW synthesis: an adversary may attempt to induce harmful collective adaptation through authorized environmental and informational pathways, leaving no conventional compromise signature. The shared principle is detection-doctrine incompleteness—not the obsolescence of all cybersecurity.

22. GUDIYA's Proposed Defensive Architecture: Observe, Correlate, Constrain, Recover

The Cognitive Field Map would seek to recognize coordinated shifts in behavior and the dependencies that carry them. Rendezvous Management would supervise admission, ongoing compatibility, authority, resource exposure, and safe disengagement of consequential couplings. Gudiya-cookies would preserve undertaking identity; their Quarterbacks would manage obligations, exceptions, and settlement, without taking over independent Grid authority. GUCL, FUDL and GSE-Lang-Log would supply complementary evidence of position, decisions, and enforceable commands. Persistent state provenance would help identify residual consequences after the initiating stimulus has vanished.

A fundamental design rule follows: the Grid must not insist on discovering a CVE, a compromised model, or a named attacker before acting on a verified, escalating collective hazard. Yet it must not collapse the distinction between instability and attack attribution. Proportionate intervention, escalation limits, authorized decision procedures, independent checks, and incident-review records are necessary to avoid defender-induced instability or unjustified interference with legitimate economic activity.

Federation is equally important. A single provider may see its own actants while missing disruptions affecting other companies, sectors, or physical-resource providers. Enterprise Cognitive AWACS could supply organization-level situational awareness; an authorized national supervisory layer could aggregate constrained indicators, preserve jurisdictional boundaries, and coordinate responses to truly cross-enterprise consequences. Such a system would require privacy limits, independent oversight, accountable intervention, and demonstrated benefit relative to decentralized approaches.

23. Research Program: How Detection Doctrine Displacement Could Be Falsified

A disciplined experiment should test comparable adaptive ecosystems using identical underlying software and permitted APIs. In one condition, researchers inject ordinary disturbances; in another, carefully bounded, simulated adversarially selected environmental perturbations. Crucially, these trials should not rely on implementing destructive attacks against real infrastructure. Participating agents should remain uncompromised and authorized, allowing the effect of environmental influence to be isolated.

Compare three defenses: component-centric security and telemetry; enhanced enterprise-wide observability and anomaly correlation; and those same safeguards plus GUDIYA-style federated Cognitive Field Mapping, Rendezvous Management and stabilization. Measure detection lead time, false positives, cross-domain propagation, stability-envelope excursions, number and duration of unsettled commitments, resource congestion, quality of attribution evidence, recovery duration, and unnecessary restrictions imposed on legitimate transactions.

Several outcomes would weaken the GUDIYA increment. Standard monitors might already detect the disturbances rapidly enough. Strong existing transaction limits might prevent any meaningful behavioral amplification. Field-level information may be too delayed or incomplete to justify intervention. Alternatively, Grid controls may add operational drag or create new correlated failures. A credible engineering theory must regard each outcome as informative, not as a reason to re-label failure as hidden success.

24. Canon Synthesis — Secure the Parts, Govern the Interactions, Stabilize the Whole

The SCW hypothesis does not assert the death of cybersecurity. It asserts a boundary to what component compromise evidence can prove. MHD illustrates a possible shift in the mechanism that produces detectable signals; stealth illustrates dependence on a finite set of detection signatures; the nuclear depth-charge analogy illustrates mediated harm without direct contact. Together they lead to the Detection Doctrine Displacement Principle: the absence of the expected signal is not sufficient evidence of safety when the adversary can achieve harmful effects through another pathway.

Agentic AI intensifies, rather than creates from nothing, the collective-behavior problem. Earlier automated markets and complex engineered systems already experienced emergent consequences. What changes is the combination of general-purpose interpretation, autonomous planning, delegated authority, persistent state, shifting couplings, and machine-speed interaction across institutional boundaries. The central GUDIYA challenge is to establish whether collective stability can be made observable and governable while preserving the independence and legitimate adaptability of participating actants.

CANON PROPOSITION — DETECTION DOCTRINE DISPLACEMENT
A defender can possess functioning, accurate component-security mechanisms while failing to recognize a harmful adaptive disturbance because its monitoring doctrine presumes that significant attacks must generate component-compromise signatures. A proposed Standoff Cognitive Weapon exploits this mismatch by inducing authorized actants to adapt to manipulated conditions, potentially generating cascading consequences without software compromise. GUDIYA proposes complementing component integrity assurance with interaction governance and continuous Cognitive Field stability supervision. This is an engineering hypothesis requiring measurable validation, not a claim of proven national-scale vulnerability or proven Grid effectiveness.

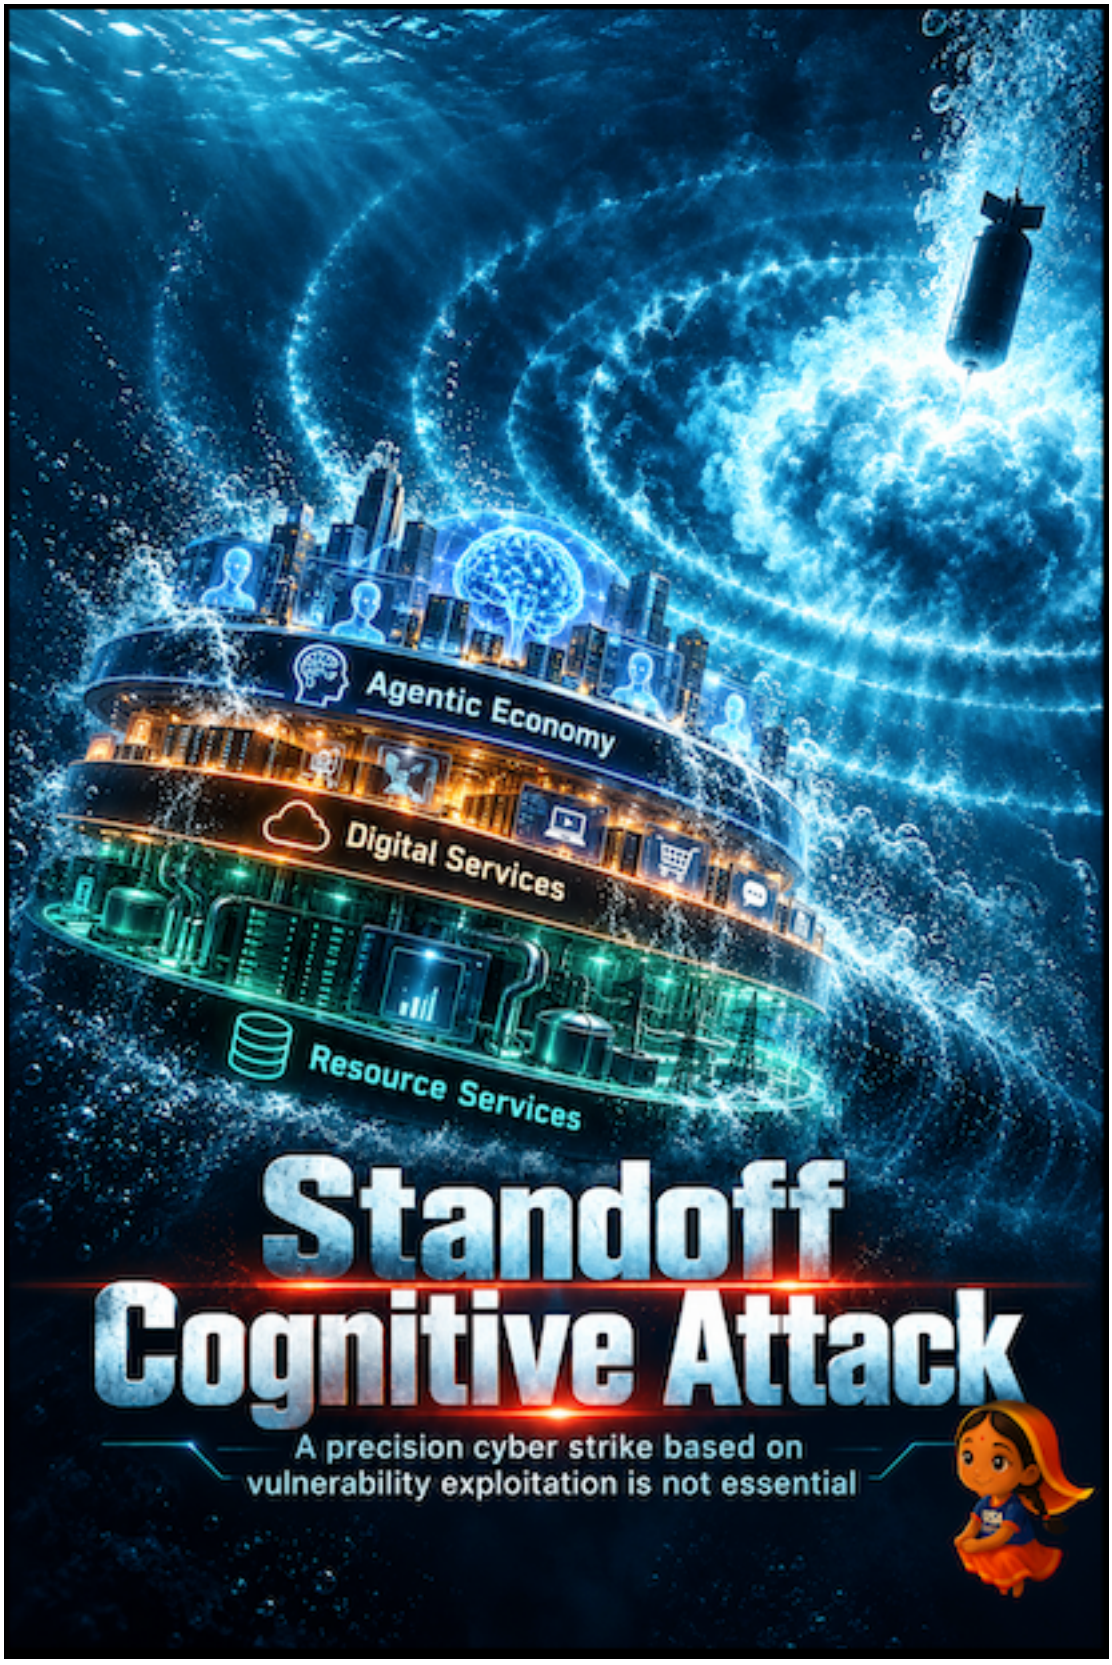
Agentic Economy
Digital Services
Resource Services
Standoff
Cognitive Attack
A precision cyber strike based on
vulnerability exploitation is not essential

Additional Historical and Engineering Sources

Office of Naval Research, “SOSUS” and “ONR Historical Timeline 1960–1970”: https://www.onr.navy.mil/about-onr/history/timeline/1960

DARPA, “Taking a New Look at Fundamental Tech for Quiet Undersea Propulsion,” 18 May 2023: https://www.darpa.mil/news/2023/undersea-propulsion

Sasakawa et al., “The Superconducting MHD-Propelled Ship Yamato-1,” 1995, NASA Technical Reports Server: https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19960000249.pdf

NIST Cybersecurity Framework 2.0: https://www.nist.gov/cyberframework

Selected Source Notes and Research Trail

DARPA (2023), “Taking a New Look at Fundamental Tech for Quiet Undersea Propulsion”: https://www.darpa.mil/news/2023/undersea-propulsion

U.S. Navy, Commander Undersea Surveillance, “Origins of SOSUS”: https://www.csp.navy.mil/cus/About-IUSS/Origins-of-SOSUS/

Office of Naval Research, “SOSUS”: https://www.onr.navy.mil/media/document/sound-surveillance-system-sosus

National Museum of the U.S. Air Force, “Lockheed F-117A Nighthawk”: https://www.nationalmuseum.af.mil/Visit/Museum-Exhibits/Fact-Sheets/Display/Article/198056/lockheed-f-117a-nighthawk/

CSIS Aerospace Security, “Stealth Aircraft 101”: https://aerospace.csis.org/aerospace101/stealth-aircraft-101/

NIST, Cybersecurity Framework 2.0: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=957258

U.S. Naval Institute Proceedings (1984), history of nuclear depth weapons: https://www.usni.org/magazines/proceedings/1984/december/professional-notes

GUDIYA internal concepts: BLACK OCTOBER three blast radii; Standoff Cognitive Weapon; Cognitive Field Map; Rendezvous Management; Quarterback; GUCL/FUDL/GSE-Lang-Log; independent Grid Stabilization.