

117-207-D NSA Guidance and GUDIYA

A Side-by-Side Comparison

Introduction — Two Different Questions About the Same Accelerating Digital World

In September 2026, the U.S. National Security Agency released “Best Practices Guide for Cyber Hygiene: Mission-focused strategies for securing networks against advanced threats.” The document responds to a changing cybersecurity environment in which artificial intelligence enables adversaries to automate and scale reconnaissance, exploitation and other stages of cyber operations. NSA consequently recommends disciplined asset management, authentication, patching, segmentation, monitoring, behavioral baselining, endpoint protection, automation, quarantine and recovery.

Read from the perspective of the GUDIYA Cognitome, the guidance is important for an additional reason. It helps expose the boundary between two engineering problems that can easily be confused:



NSA asks, principally:

- How do we prevent adversaries from exploiting increasingly fast and automated digital systems?


GUDIYA asks:

- How do we keep increasingly autonomous, interconnected and machine-speed cognitive systems dynamically stable—even when nobody is attacking them?
- Both problems matter.
- But they are not the same problem.

1. NSA Begins With an Adversary

- The NSA guidance describes a threat environment in which AI can accelerate and scale computer-network exploitation. Its recommendations address familiar but consequential weaknesses including unpatched systems, weak authentication, insecure configurations, inadequate segmentation and insufficient monitoring.

- Its conceptual world can therefore be represented approximately as:



- AI changes the velocity and scale of this process.
- Hence:

$AI + \text{Adversary} \rightarrow \text{FasterAttack}$

- The defensive response is correspondingly stronger cybersecurity.

2. GUDIYA Begins With Dynamics

- GUDIYA starts from a different systems-science hypothesis.
- Consider four legitimate actants: A,B,C,D
- Every actant can be:
 - authenticated,
 - authorized,
 - uncompromised,
 - properly patched,
 - correctly configured,
 - pursuing an approved objective.
- There may be no adversary whatsoever.
- Yet connect them:

$$A \leftrightarrow B \leftrightarrow C \leftrightarrow D$$
- Add machine-speed cognition:

$$\text{Velocity} \uparrow$$
- Add recursive delegation:

$$\text{Recursion} \uparrow$$
- Add feedback:

$$\text{Gain} \uparrow$$
- Add increasingly dense rendezvous:

$$\text{RendezvousDensity} \uparrow$$

and the system can potentially acquire dynamics not meaningfully attributable to any single actant:

 - Oscillation
 - Amplification
 - PersistentObjectives
 - SyntheticGroups
 - Cascades
 - Emergence
- Thus GUDIYA's problem formulation becomes:

LegitimateActants + Coupling + Feedback + Recursion + MachineSpeed → PossibleInstability

- No intrusion is necessary.

3. The Most Important Test Case

Imagine a future enterprise in which NSA's recommendations have been implemented superbly.

- Every asset is inventoried.
- Every identity is authenticated.
- Privileges follow least privilege.
- Software is patched.
- Network segmentation is functioning.
- Endpoint detection is operating.
- Behavioral monitoring is active.
- Logs are collected.
- Recovery procedures have been tested.

- There is:

NoIntrusion
NoMalware
NoCredentialTheft
NoUnauthorizedAccess

- Now imagine 500 authorized AI actants interacting continuously.
- They begin recursively delegating work, creating temporary groups, sharing state, responding to one another, invoking tools and changing their behavior in response to the environment they collectively create.
- Suddenly:

CouplingGain ↑
Recursion ↑
CognitiveDensity ↑
CognitiveVelocity ↑
StabilityMargin ↓

- The enterprise begins moving:

$S0 \rightarrow S1 \rightarrow S2 \rightarrow S3$

- Nothing about this hypothetical condition implies that NSA's cybersecurity controls failed.
- They may be functioning exactly as intended.
- The failure belongs to another class:

Dynamic Cognitive Instability

- This is the space GUDIYA proposes to address.

4. The Side-by-Side View

Dimension	NSA Cyber Hygiene Guidance	GUDIYA Cognitive Stability
Primary discipline	Cybersecurity	Cognitive Stability Engineering
Primary concern	Cyber exploitation and compromise	Dynamic instability of machine-speed cognition
Threat origin	Predominantly adversarial	May be entirely non-adversarial
AI's significance	Accelerates/scales cyber exploitation	Creates autonomous cognition and machine-speed interaction
Protected object	Networks, systems, identities, data	Cognitive Field and its actants/interactions
Fundamental question	"Has something been compromised?"	"Are the dynamics becoming unstable?"
Asset identity	Asset inventory and identity management	KYA — Know Your Actant
Access	Authentication / least privilege	Authority and Cognitive Envelopes
Topology	Network topology/data flows	Effective Cognitive/Causal topology
Segmentation	Network segmentation	Cognitive separation and firebreaks
Observability	Logs, traffic, endpoint telemetry	GCR/Grid telemetry and Cognitive Field observability

Dimension	NSA Cyber Hygiene Guidance	GUDIYA Cognitive Stability
Baseline	Normal network/system behavior	Normal dynamic cognitive behavior
Anomaly	Possible intrusion or malicious behavior	Possible instability or trajectory deviation
Interaction model	Connections/data flows	Rendezvous and dynamic coupling
Implicit coupling	Not a focus of this guide	Core concern
Coupling gain	Not addressed	Core stability variable
Cognitive Velocity	Not addressed	Core GUDIYA concept
Cognitive orientation	Not addressed	CCS / Cognitive Attitude / external reference
Intent trajectory	Not addressed	Intent-Group Plan
Operating boundary	Security policies / permissions	Cognitive Envelope
Recursion instability	Not addressed	Core concern
Emergence	Not a focus of this guide	Fundamental concern
Persistent objectives	Not addressed	Cognitive Field phenomenon to observe
Automation	Defensive force multiplier	Useful but itself a potential dynamic actant
Automated quarantine	Recommended in appropriate security workflows	Dynamic quarantine based on instability
Isolation trigger	Security anomaly / suspected compromise	Dynamic instability may suffice
Response philosophy	Defense-in-depth and incident response	Minimum sufficient stabilization
Escalation	Tiered cybersecurity maturity/response capabilities	Observe → Advise → Shape → Constrain → Separate → Quarantine → SCRAM
Ultimate concern	Prevent/contain compromise	Prevent S4 systemic collapse
Recovery	Restore secure operations	Restore stable cognition and progressively return autonomy

The comparison reveals considerable architectural rhyme, while the underlying failure models remain different.

5. Inventory Becomes KYA

- One of NSA's foundational recommendations is knowing what exists within the environment.
- This is elementary but profound:
CannotProtect → WhatYouCannotIdentify
- GUDIYA independently arrives at an analogous requirement for the Cognitive Field:

KYA — Know Your Actant

- But an actant inventory requires richer information than a conventional IT asset inventory.
- GUDIYA needs to know not merely:

- What is this thing?
- but potentially:
 - What cognition can it perform?
 - What authority does it possess?
 - Which tools can it invoke?
 - What Cognitive Envelope applies?
 - What other actants can it rendezvous with?
 - What recursion or delegation can it initiate?
 - What is its RM rating?
 - What is its SCRAM behavior?
- Thus:

AssetInventory → KYA

is not replacement. It is an extension necessitated by cognition.

6. Network Topology Becomes Cognitive Topology

- NSA emphasizes understanding assets, connections and data flows and maintaining useful topology information.
- Cybersecurity therefore wants to understand:

WhoCanConnectToWhom?

- GUDIYA needs another graph:

WhoCanInfluenceWhom?

- The distinction can be enormous.
- Suppose:

$A \nrightarrow B$

according to the declared network architecture.

- But both interact through:

$A \rightarrow \text{Database} \rightarrow B$

or:

$A \rightarrow \text{Queue} \rightarrow B$

or:

$A \rightarrow \text{SharedMemory} \rightarrow B$

- Then:

NetworkTopology $\neq$ CognitiveCausalTopology

- This is GUDIYA's Implicit Coupling problem.
- The architecture tells GUDIYA who is supposed to interact.
- The Cognitive Field must reveal who is actually influencing whom.

7. Behavioral Baselines Become Dynamic Baselines

- NSA recommends establishing baselines and monitoring for anomalous behavior.
- Again, GUDIYA agrees with the operational philosophy.
- But the observables change.
- NSA might observe:

Traffic
Accounts
Applications
Connections

Processes
Privileges

- GUDIYA additionally wants:
CognitiveVelocity
CognitiveAttitude
CouplingGain
RendezvousDensity
RecursionRate
IntentDrift
CognitiveDensity
CognitiveCombustibility
PersistentObjectiveStrength
StabilityMargin
- This produces a useful distinction:

CybersecurityObservability: WhatIsHappening?

versus:

CognitiveStabilityObservability: WhatAreTheDynamicsBecoming?
- The second question is inherently predictive.

8. Segmentation Becomes Cognitive Separation

- NSA strongly values segmentation because it can limit lateral movement and contain compromise.
- GUDIYA inherits the principle but changes the reason for invoking it.
- Cybersecurity may separate:
AB
because A is compromised.
- GUDIYA may separate:
AB
even when:
Secure(A) = TRUE
and:
Secure(B) = TRUE
because:

Stable(A ↔ B) = FALSE
- That is a profound distinction.
- The interaction itself can become the object requiring intervention.
- Hence:
Cybersecurity → SecureTheActant
while GUDIYA adds:
CognitiveStability →

StabilizeTheInteraction

9. NSA Quarantine and GUDIYA Dynamic Quarantine

- NSA includes automated response and quarantine among progressively mature defensive capabilities.
- This provides another striking parallel.
- But compare the triggers.

- Cybersecurity
SecurityAnomaly → PossibleCompromise → Quarantine
- GUDIYA
DynamicAnomaly → PossibleInstability → Quarantine
- An actant might therefore be quarantined by GUDIYA while remaining:
CyberStatus = GREEN
because:
StabilityStatus = RED
- GUDIYA quarantines not merely bad actors.
- She may quarantine bad dynamics among good actors.

10. Automation Creates the “Defend the Defenders” Problem

- NSA encourages increasing defensive automation, including orchestration and automated response.
- That is entirely understandable.
- Machine-speed attacks increasingly invite machine-speed defense.
- But GUDIYA asks one additional question:

WhoStabilizesTheAutomatedDefender?

- Imagine:
DetectionAgent → SOAR → EDR → IdentityAgent → FirewallAgent
- These defensive components observe the enterprise and alter it.
- Their interventions change subsequent telemetry.
- That telemetry changes their subsequent interventions.
- Therefore:

Defender → System

and:

System → Defender

produce:

Defender ↔ System

- The supervisory architecture has entered the dynamics.
- This is the GUDIYA Supervisory Multi-Agent Stability Principle:
- A system designed to supervise another system does not acquire immunity from instability merely because its purpose is defensive.
- Thus emerges GUDIYA's First Principle:

DEFEND THE DEFENDERS

- Firefighters need buddy systems.
- Stabilizers need stabilization.

11. Static Configuration Versus Dynamic Stability

- Cybersecurity necessarily contains substantial preconfiguration:
Identity
Permissions
Segmentation
Allowlisting
Policies
PatchLevels

- These controls can subsequently be monitored and dynamically enforced.
- Cognitive Stability begins from a somewhat different premise:
- A coupling that is safe now may become unstable later without either participant changing identity or violating policy.
- Therefore:

$\text{Safe}(A, B, t_0)$

does not guarantee:

$\text{Safe}(A, B, t_1)$

because:

Context

Velocity

Density

Gain

Recursion

Topology

may have changed.

- Hence GUDIYA's stabilization problem is dynamic by default.

12. The Missing Variable Is Emergence

- Perhaps the largest distinction concerns emergence.
- Cybersecurity generally looks for deviations that indicate something undesirable has happened:

Intrusion

Malware

CredentialAbuse

UnauthorizedAction

- GUDIYA must additionally observe something much stranger:

$\text{SomethingNewCanAriseFromInteractionsAmongLegitimateActants}$

- No actant individually needs to contain the resulting behavior.
- The behavior may belong to:

The System

rather than:

An Agent

- This is why GUDIYA treats agentic AI as a Complex Adaptive Systems problem rather than merely a cybersecurity problem.

13. Cybersecurity Can Be Perfect and Stability Can Still Fail

- This deserves canonical status:

$\text{CybersecuritySuccess} \Rightarrow \text{CognitiveStability}$

- Similarly:

$\text{CognitiveStability} \Rightarrow \text{CybersecuritySuccess}$

- A perfectly stable system can be hacked.
- A perfectly secure system can become dynamically unstable.
- Therefore:


$\text{Security} \perp \text{Stability}$

- They are partially orthogonal engineering dimensions.

- A future AI infrastructure needs both.

14. From Defense-in-Depth to Stability-in-Depth

- NSA advocates layered defense rather than dependence upon one control.
- GUDIYA can learn directly from this philosophy.
- Cognitive Stability should likewise never depend upon one brilliant algorithm.
- Instead:



- This might appropriately be called:

Stability-in-Depth
- No single defensive mechanism carries the entire burden.

15. Two Escalation Philosophies

- NSA's guide presents progressively mature cybersecurity practices, from immediate foundational protections toward greater automation, orchestration, adversarial emulation and validation.
- GUDIYA requires a runtime escalation structure:

Observe → Advise → Shape → Constrain → Separate → Quarantine → SCRAM

with the governing principle:

MinimumSufficientIntervention

- GUDIYA does not want maximum control.
- She wants:

MaximumSafeAutonomy
- The Grid increases authority only as stability margin deteriorates.

16. Cognitive SCRAM Has No Direct NSA Equivalent

- NSA discusses containment, automated response and recovery.

- GUDIYA's extreme failure model creates an additional concept.
- If:

StabilityMargin → Critical

and:

S4 → Imminent

the Grid's objective changes from:

PreserveProductiveCognition

to:

PreventS4

- Then comes:

COGNITIVE SCRAM

- Cognitive SCRAM is not another security quarantine.
- It is GUDIYA's proposed pre-engineered emergency transition of affected actants toward certified bounded safe states.
- The Grid stops trying to optimize cognition.
- It arrests the degrees of freedom through which catastrophic cognitive instability could continue propagating.

17. The Maturity Continuum

- The comparison reinforces an earlier GUDIYA proposition:

$\text{Cybersecurity} \subset \text{CognitiveSecurity} \subset \text{CognitiveStability}$

- This should not be interpreted to mean that Cognitive Stability replaces cybersecurity.
- The circles represent widening problem domains.

Cybersecurity asks:

Is an actant, network or resource secure against unauthorized or malicious activity?

Cognitive Security asks:

Is cognition being manipulated, leaked, poisoned, redirected or exercised outside legitimate authority?

Cognitive Stability asks:

Regardless of whether every participant is legitimate, are the collective dynamics remaining inside acceptable stability bounds?

The outer circle therefore must coexist with the inner circles.

18. A More Complete Future Enterprise Defense Stack

- The NSA document represents an increasingly sophisticated cybersecurity defense.
- GUDIYA proposes adding another infrastructure layer:

Cybersecurity

protecting:

Assets + Identity + Networks + Data



Cognitive Security

protecting:

Intent + Context + Authority + Cognition



Cognitive Stability

protecting:

Interactions + Couplings + Dynamics + Emergence

- Together:

Secure + CognitivelySecure + DynamicallyStable

becomes the desired condition for machine-speed enterprise cognition.

19. What GUDIYA Should Learn From NSA

- Cybersecurity has learned through decades of experience that serious protection requires:
 - Inventory. Identity. Baselines. Continuous telemetry. Defense-in-depth. Segmentation. Automated response. Isolation. Exercises. Recovery.
 - Those principles translate remarkably well into Cognitive Stability Engineering.
 - The translation might look like:

Mature Cybersecurity Lesson	GUDIYA Translation
Know your assets	Know Your Actant (KYA)
Know your topology	Map the Cognitive Field
Establish baselines	Establish dynamic stability baselines
Continuously monitor	Continuous Grid observability
Least privilege	Bounded actant authority
Segment networks	Bound Cognitive Spaces / create firebreaks
Detect anomalies	Detect dynamic instability
Automated response	GRE stabilization intervention
Quarantine	Dynamic quarantine
Defense-in-depth	Stability-in-Depth
Incident exercises	Instability Preparedness Drills
Test recovery	Stable-state recovery drills
Assume breach	Assume emergence
Protect defenders	Defend the Defenders

- The final pair is particularly important.
- Cybersecurity matured from:
 - Prevent every intrusion.
- toward:
 - Assume breach and engineer resilience.
- Cognitive Stability may eventually mature similarly:
 - Do not assume every emergent behavior can be prevented. Assume emergence—and engineer the Cognitive Field so emergence does not automatically become instability.
- That may be one of the most valuable lessons GUDIYA can inherit from NSA's discipline.

20. What NSA Can Eventually Learn From Stability Engineering

- The relationship can eventually become reciprocal.
- As cybersecurity systems themselves become increasingly agentic:

HumanDefender → AIObserver → AIDetector → AIResponder → MultiAgentDefensiveSystem

cybersecurity infrastructure itself becomes an HCAS.

At that point:

CybersecuritySystem

becomes:

ObjectRequiringCognitiveStabilization

The system defending the enterprise may itself need:

- Cognitive Envelopes,
- external reference,
- Rendezvous Management,
- coupling monitoring,
- recursion limits,
- Cognitive Attitude,
- Grid observability,
- dynamic quarantine,
- escalation discipline,
- SCRAM compliance.

Thus GUDIYA does not sit in opposition to NSA-style cybersecurity.
Eventually it may sit around it.

21. The Boundary of Responsibility

This gives us a clean division.

NSA-style cybersecurity:

ProtectTheDigitalEnvironmentFromThreatActors

GUDIYA-style Cognitive Stability:

ProtectTheCognitiveFieldFromUnstableDynamics

- There will naturally be overlap.
- A cyberattack may induce Cognitive Instability.
- A Cognitive Instability may expose new cybersecurity vulnerabilities.
- But neither discipline subsumes the other.
- They defend against different classes of failure.

22. The Side-by-Side Thought Experiment

Imagine two warning lights on the future enterprise dashboard.

- The first says:

CYBERSECURITY: GREEN

- No detected intrusion.
- No compromised credentials.
- No unauthorized access.
- No malware.
- No known exploitation.
- Beside it:

COGNITIVE STABILITY: RED

- Coupling gain rising.
- Rendezvous density accelerating.
- Recursion expanding.
- Cognitive orientation degrading.
- Cognitive Density Amplification underway.
- Stability margin collapsing.
- That dashboard summarizes the entire distinction.
- The system can be secure and unstable simultaneously.

And, conversely:

CYBERSECURITY:	RED
COGNITIVE SECURITY:	GREEN

- is also possible.
- The two instruments measure different dimensions of system health.

23. The Larger Historical Lesson

- NSA's document also reminds GUDIYA of something larger.
- Cybersecurity did not emerge fully formed.
- As computing became interconnected, society gradually discovered new failure modes.
- Those discoveries produced:

Firewalls
IDS/IPS
SIEM
ZeroTrust
EDR
SOAR
ContinuousMonitoring

- Each represented adaptation to a changing systems environment.
- Agentic AI may be initiating another transition:

StandaloneComputing → NetworkedComputing → CloudComputing → APIEconomy
→ AgenticCognitiveField

- The new environment does not invalidate cybersecurity.
- It exposes another engineering problem layered above it.

Conclusion — NSA Secures the Network; GUDIYA Stabilizes the Cognitive Field

- NSA's September 2026 Cyber Hygiene guidance represents disciplined, mature thinking about protecting digital infrastructure against increasingly sophisticated and AI-enhanced threats.
 - GUDIYA should learn from it.
 - Inventory matters.
 - Identity matters.
 - Segmentation matters.
 - Continuous observability matters.
 - Baselines matter.

- Automation matters.
 - Quarantine matters.
 - Exercises matter.
 - Recovery matters.
 - Defense-in-depth matters.
- But GUDIYA begins where the cybersecurity problem statement eventually ends.
- Imagine the adversary disappears.
- Imagine every vulnerability has been patched.
- Imagine every actant is authenticated.
- Imagine every interaction is authorized.
- Imagine every cybersecurity dashboard reads:
- GREEN
- GUDIYA still asks:
 - Is the Cognitive Field dynamically stable?
 - Because legitimate actants can couple.
 - Couplings can form feedback.
 - Feedback can amplify.
 - Recursion can accelerate.
 - Interactions can produce emergence.
 - Emergence can acquire persistence.
 - Cascades can propagate.
 - And a system composed entirely of individually functioning components can nevertheless become unstable as a whole.
- That leads to the central lesson of this comparison:

SecurityAsks: IsSomeoneDoingSomethingTheyShouldNot?
- while:

StabilityAsks: IsTheSystemBecomingSomethingItShouldNot?
- The NSA guidance strengthens the defenses of the digital infrastructure upon which future cognition will run.
- GUDIYA proposes another layer around that infrastructure:

Cybersecurity → CognitiveSecurity → CognitiveStability
- The objective is not to choose between them.
- The future enterprise will require all three.
- NSA teaches us how to defend increasingly intelligent networks from increasingly intelligent adversaries.
- GUDIYA asks the next systems-engineering question: who stabilizes the intelligent network when the network itself becomes a cognitive system?

