Cyber Struggle
SPECIAL FORCES MINDSET IN CYBER WORLD

G.R.A.S.P

Beyond Tool and Technology Validation

S46 GRASP Whitepaper

A Multidimensional Framework

Measuring Human Readiness in Incident Response Execution

> My mind is my best weapon.

February 24th
— 2026

version 1.0

Document History

1.0 | Mart 2026 | Cyber Struggle | First Publish

Table of Content

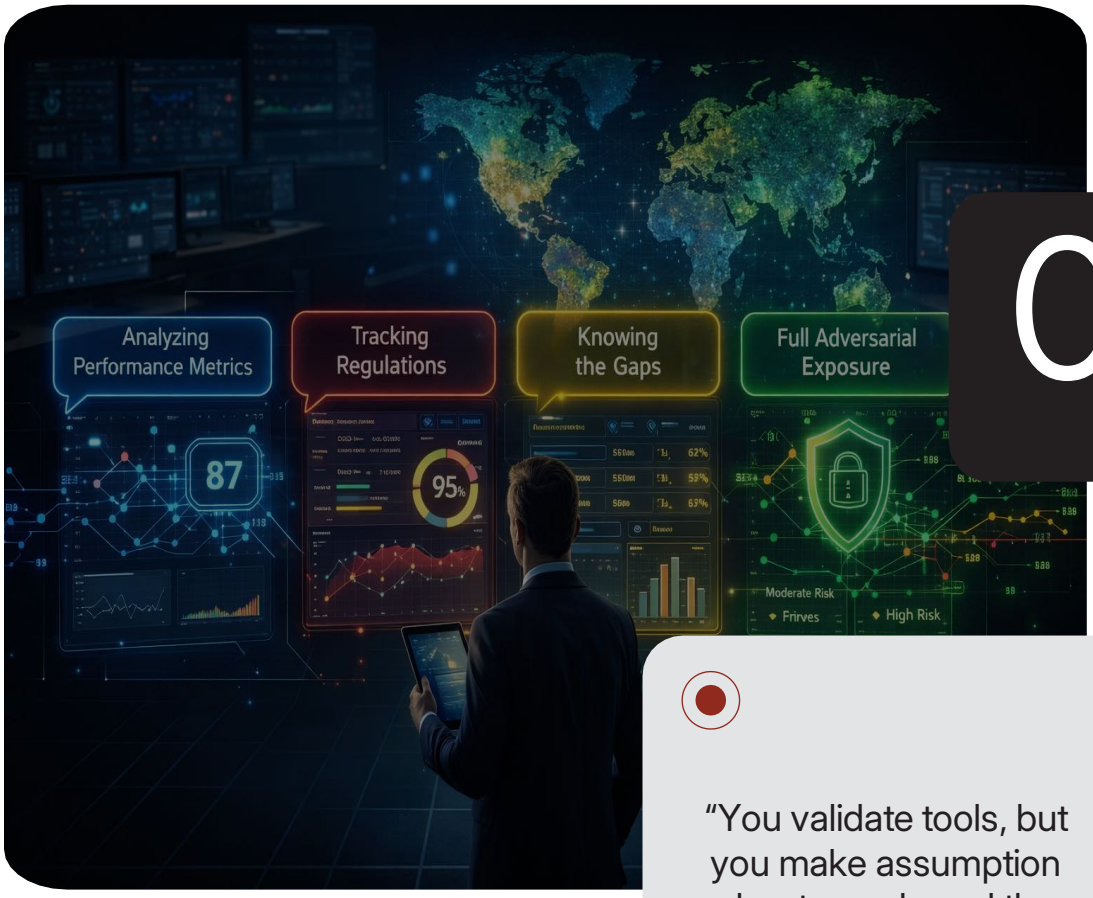

01

"You validate tools, but you make assumption about people and the team!"

Executive Summary

Executive Summary

Organizations have grown significantly more mature in validating security tools, controls, and detection logic. Yet when a real incident unfolds, they rarely apply the same rigor to measuring whether people can investigate effectively, whether analysts can build context under pressure and sustain decision quality when it matters most.

This gap is the readiness blind spot.

Incident readiness today is largely assessed through proxy indicators: certifications, completed training, years of experience, managerial confidence. These signals may reflect capacity, but they cannot measure where performance breaks down under operational pressure. As a result, organizations typically discover the limits of their team's readiness during an actual incident. **What organizations cannot measure, they cannot trust under pressure!**

GRASP is a readiness validation system designed to close this gap. It picks up where tool validation ends. Once signals are generated, GRASP makes visible what actually happens at the human and process layer. The objective is not to measure what analysts know. It is to observe how investigation quality shifts under incomplete information, conflicting signals, time pressure, and crisis communication.

The core thesis of this whitepaper is this: incident readiness is not the sum of stored knowledge. Readiness is the ability to build context, generate actionable hypotheses, update mental models against conflicting evidence, escalate at the right moment, and remain coherent under pressure. Because under high-pressure, high-ambiguity conditions, analysts do not rise to their best. They fall to their level of preparation. This is why readiness can only be validated in environments that resemble real operations: unpredictable, context-incomplete, and behaviorally observable.

This framework draws on two converging bodies of research: the neuroscience of prediction error and active model updating, and the educational science of situated and experiential learning. People do not become more capable by storing knowledge alone. They develop by taking action in live context, receiving feedback, and revising their mental models. GRASP brings this process into the operational security environment. It makes an analyst's investigation flow, communication under pressure, and decision architecture visible, then converts those observations into development outcomes.

GRASP is not positioned as a training platform. It does not replace BAS, tabletop exercises, or cyber range investments. It completes them, with the one performance layer most organizations have never been able to directly measure: the readiness of their people, their processes, and ultimately their teams.

You're accountable for readiness you cannot measure

02

Readiness Blind Spot

Up to 64% analysts facing burn-out

Readiness Blind Spot

Modern security teams are getting increasingly effective at testing what is easiest to measure. Detection rules can be evaluated, control coverage can be mapped, playbooks can be written, and processes can be audited. All of these are essential components of a mature security program.

However, on their own, they do not answer a fundamental operational question: when conditions are uncertain and pressure is high, how well does the response team actually perform?

In most organizations, analyst readiness is not directly or systematically observed. Instead, proxy indicators are used: certifications, years of experience, completed training, participation in tabletop exercises, and managerial confidence. These inputs can provide useful signals, but they do not reflect the quality of investigations under real operational conditions. As a result, organizations often only discover the limits of their teams during an actual incident.

This blind spot becomes most visible in real cases. One analyst rapidly builds a working mental model from the available data, develops useful hypotheses, and progresses decisively. Another, with similar technical knowledge, may collect excessive data, pursue weak investigative paths, delay decision-making, or mistime escalation. The same incident can be contained in forty minutes by one analyst, while it may take two hours with another.

This is not a rare exception. In many SOC and incident response environments—especially in MSSP structures where multiple analysts, shifts, and client contexts intersect—this is a recurring operating pattern. The issue is not a lack of effort. Many organizations have comprehensive tooling, mature processes, and extensive training programs. The real problem is the absence of a reliable way to observe the quality of human response under realistic operational pressure.

For this reason, the readiness blind spot is not just a training problem; it is also a measurement problem. What cannot be measured cannot be managed—and what is managed through proxy signals often breaks down under real pressure.

74% of Team Members Quit

Because of

Stress

03

The Limits of Current Readiness Methods

The scale of this gap is supported by both industry data and observational evidence. According to IBM's 2025 Cost of a Data Breach report, the average breach lifecycle still stands at 241 days; organizations using AI and automation are able to reduce this by roughly 100 days. A significant portion of this difference comes not just from technology, but from people making better decisions.

Industry signals are also increasingly pointing to human performance under pressure as a distinct problem space. According to Cynet's 2023 CISO survey, 94% of CISOs report experiencing stress at work, and 64% say this stress directly weakens their organization's ability to defend itself. The 2025 SANS & GIAC Cybersecurity Workforce Research Report further shows that hiring profiles in cybersecurity have shifted fundamentally: what used to be a 70% technical / 30% attitude balance has largely reversed, with adaptability and psychological resilience becoming primary differentiators. The same report indicates that 55% of organizations consider on-the-job experience to be the most effective method of development.

Cyber Struggle's Ranger program, conducted with over 400 participants, provides a data foundation that allows this gap to be directly observed. Across 23 cohorts, participants were evaluated on technical problem-solving speed, quality of critical thinking, and psychological resilience. Three findings stand out: participants trained through a combination of situated and experiential learning demonstrated at least 60% higher functionality and resilience—compared to their baseline—when dealing with previously unseen incidents under operational pressure. Operational consistency increased by up to five times. Perhaps most notably, participants with less prior domain knowledge were able to produce more effective and better-balanced solutions to first-time problems than those with greater knowledge, given sufficient methodological exposure. While these observations, derived from a development program without a control group, do not establish causality, they do demonstrate that human performance can be measured, improved, and that inconsistency is not inevitable.

This final observation reinforces GRASP's core premise: functional readiness is not determined by the volume of stored knowledge, but by the ability to construct meaning, test assumptions, and adapt under uncertainty.

Most organizations are not ignoring readiness altogether. They invest in BAS platforms, tabletop exercises, labs, cyber range environments, certifications, and playbook development. The issue is not a lack of effort, but the fact that these methods validate only a portion of the capacity required for real incident response.

Breach and Attack Simulation tools are effective at validating whether security controls and detection logic are functioning as expected. However, they do not show how an analyst interprets a signal once an alert is generated, how they build context, or how they drive the investigation forward. Tabletop exercises are valuable for coordination and decision-making frameworks, but participants typically know what is being tested, and the pressure rarely reflects operational reality. Tabletop exercises test discussion, not behavior. Cyber ranges, labs, and CTF-style environments are effective for developing isolated technical know-how, but uncertainty, disruption, crisis communication, and evolving context are often secondary.

The most operational question remains unanswered: when information is incomplete, pressure is rising, signals are conflicting, and time does not wait for clarity; can teams maintain the quality of their investigations?

Same Alert. 40 minutes or 2 hours – You don't know

> What surprised me most wasn't only technical improvement; it was psychological shift. Analysts who used to hesitate or over-analyze became more decisive. Get better with incomplete information instead of waiting certainty. GRASP forces us repeated exposure to this pressure.
>
> **SOC Manager - Telecommunication**

why

“Know but Freeze” and Analysts Performance

04

54% of incidents are caused by lack of incident response readiness, not tool failure.

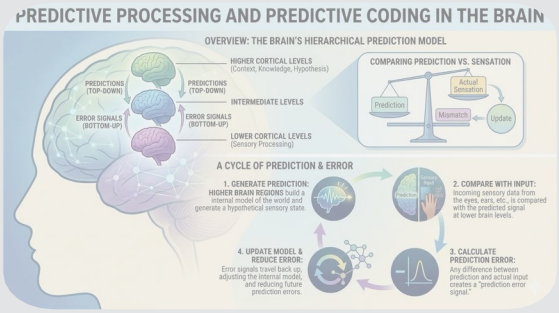

Human perception and decision-making are increasingly explained through the framework of predictive processing. In this view, the mind is not a passive system that processes incoming data, but an active modeling system that continuously generates predictions about possible states of the world.

Within this model:

- The mind continuously generates expectations (predictions)
- Incoming signals are compared against these expectations
- Mismatches between expectations and incoming input emerge as prediction error
- This mismatch drives the updating of the internal model

Cognition, therefore, is not driven by data alone, but by the interaction between expectation and evidence. The quality of performance depends not on the complete elimination of prediction error, but on how effectively this mismatch is evaluated under conditions of uncertainty (Clark, 2013; Friston, 2010).

This framework provides a strong lens for understanding incident response.

Analysts do not process data passively; they construct and continuously update working hypotheses and mental models about what is happening in the environment.

- Alerts and telemetry → incoming signals
- Working hypotheses → predictions
- Conflicting evidence → prediction error
- Investigation → the process of testing hypotheses and updating the model

From this perspective, incident response is not a data processing problem, but a continuous cycle of hypothesis generation, testing, and revision under uncertainty.

Incident response is not merely about processing data; it is the continuous construction and updating of a mental model under uncertainty.

Analysts do not encounter incidents as passive recorders of information. They interpret signals through prior experience, learned patterns, and working hypotheses about the most plausible scenario.

High performance is not about rejecting conflicting evidence, but about incorporating it into the model updating process. Strong analysts, when confronted with new evidence, update, refine, or restructure their models.

Poor performance, by contrast, typically emerges at two extremes:

- Overcommitment to an initial hypothesis
- Overreaction to each new signal, leading to loss of directional coherence

In both cases, the outcome is the same: progress stalls and effort shifts toward unproductive data accumulation.

For this reason, incident response is inherently iterative.

This leads to a critical implication:

If the environment is overly clean, overly guided, or overly predictable, analysts are not forced to meaningfully challenge or update their models. Prediction error remains limited, learning stays superficial, and acquired knowledge often fails under real uncertainty.

This principle sits at the core of GRASP's design philosophy:

real adaptation emerges only under real friction and uncertainty.

From this perspective, readiness is not the accumulation of stored knowledge;

it is the ability to construct, test, and continuously update meaning under uncertainty.

True readiness is repeatable, that is

resilience

05

> For the first time, we can see how our analysts behave in real incidents; how they think, where they hesitate, and where performance breakdown
>
> **CISO – Energy Sector**

Relation with Situated and Experiential Growth

This cognitive model intersects with two key frameworks that explain how human performance develops: situated learning and experiential learning (Kolb, 1984; Lave & Wenger, 1991).

From a situated learning perspective, knowledge is not a context-independent asset. People do not simply learn what to do; they internalize how that knowledge is applied under specific conditions; through particular signals, within certain priorities, and under social or operational pressure. This is critical in a domain such as incident response, where performance depends not only on tool proficiency, but on the ability to act appropriately and in a timely manner within a dynamic context (Lave & Wenger, 1991).

Experiential learning explains how this capability develops. Operational readiness does not emerge from exposure alone; it requires a cycle of action, feedback, reflection, adaptation, and reapplication. An analyst encounters a situation, forms a working hypothesis, takes action, observes outcomes, identifies where the model breaks down, and constructs a more refined model in the next iteration (Kolb, 1984).

Taken together, these two approaches lead to a critical conclusion: functional readiness is not defined by the accumulation of knowledge, but by the ability to activate that knowledge in context and continuously update it through learning from error. The friction generated by prediction error becomes the driving force of the experiential learning cycle, while situated context ensures that this development transfers to real operational environments (Kolb, 1984; Lave & Wenger, 1991).

For GRASP, this combination forms a strategic foundation. The product is not designed to deliver information or present scenarios. Its purpose is to place the analyst in a realistic and dynamic context, generate meaningful prediction error, make investigation behaviors observable, and progressively develop those behaviors toward greater operational consistency.

GRASP's objective is not to provide content, but to develop operational readiness through realistic friction, uncertainty, and feedback loops: to make investigation behavior visible and to make performance more consistent over time.

You validate tools continuously But you assume people

> Managing multiple analysts across shifts, we always knew there was variance, but we couldn't quantify it. GRASP made that visible immediately. Same type of incident, completely different investigation path, crisis communication approach and resilience.
>
> **SOC Manager - MSSP**

06

Designing Effective Readiness Exposure

Readiness is not a single-dimensional capability. Real incident response performance is the result of multiple interacting factors: technical knowledge, contextual reasoning, hypothesis quality, cognitive flexibility under stress, emotional regulation, decision discipline, and the ability to act consistently under operational pressure.

For this reason, an effective readiness validation model must be designed not only to assess technical correctness or process adherence, but to make this multidimensional structure of human performance observable.

In this context, readiness validation must focus on directly observing how analysts apply these capabilities under realistic conditions. This requires five key conditions.

Unpredictability. Real incidents do not follow a schedule. Analysts are not presented with fully defined scenarios or clearly bounded problem statements at the outset. A validation environment must reflect this uncertainty.

Incomplete and evolving context. Signals do not arrive as clean, self-explanatory packages; they are fragmented, sometimes conflicting, and only gain meaning as the investigation progresses. When context is provided upfront, the analyst's core task—constructing meaning—is no longer observable.

Operational pressure. Time pressure, interruptions, crisis communication, and management visibility directly shape how analysts think and make decisions. When these elements are removed, what remains is a safe exercise—not readiness validation.

Behavioral observability. Not only the outcome, but the investigation process leading to that outcome must be observable: when the initial hypothesis was formed, which directions were explored, where the analyst became stuck, how quickly they adapted to conflicting evidence, and how escalation decisions were made.

Developmental value. Validation should not end with a score. The signals generated must translate into coaching, standardization, and targeted development actions.

Turn readiness from assumptions into evidence

07

How GRASP Operates This Framework

> With increasing regulatory pressure, especially around operational readiness, for the first time, incident readiness becomes auditable with GRASP for us.
>
> **CISO – Finance**

GRASP is designed as a response to this measurement problem.

The starting point is clear: what is missing in the cybersecurity validation space is not another way to measure whether systems detect attacks, but a layer that makes human performance visible after an incident begins. GRASP focuses on the human dimension of incident response: how analysts investigate, adapt, prioritize, and communicate as uncertainty and pressure increase.

To achieve this, GRASP examines the investigation process itself. It focuses on where the analyst starts, how quickly they arrive at the first meaningful hypothesis, which directions they explore, how they respond to conflicting evidence, when they make escalation decisions, and how they maintain communication under pressure. Stressor-driven interruptions and crisis communication dynamics are incorporated into the evaluation design, ensuring that not only technical skills, but also cognitive regulation, prioritization, and communication quality become observable.

AI-powered behavioral analysis transforms this investigation process into meaningful readiness signals. The goal is not to label the analyst as superficially “good” or “bad,” but to identify where response quality breaks down: in hypothesis quality, unnecessary branching, adaptation to conflicting evidence, communication, or escalation timing.

This connection is critical. Prediction error is necessary for learning, but it is not sufficient on its own. If friction is not made visible and linked to structured feedback, the analyst is merely challenged without meaningful development. What GRASP does is convert this friction into measurable signals and connect them to targeted development actions. In doing so, it operationalizes the model-updating logic from neuroscience and the experiential learning cycle from learning sciences within real-world security operations (Clark, 2013; Friston, 2010; Kolb, 1984; Lave & Wenger, 1991).

Equally important is the developmental logic underlying this system. During the initial stages of incident exposure and hypothesis formation, GRASP operates through a situated learning lens, making visible how analysts interpret context, signals, and operational conditions. As the process evolves and conflicting evidence emerges, prediction error comes into play and the focus shifts toward the experiential learning cycle. At this stage, the analyst tests their working hypothesis, receives feedback, updates their mental model, and develops a more functional approach in the next iteration. In this way, GRASP does not merely provide incident exposure; it operationalizes predictive processing within a real-world context, systematically developing the analyst’s ability to construct, test, and adapt models under uncertainty.

The Need

Situated Context + Experiential Loop + Observable Prediction Error

GRASP

Unpredictable incident, stressors, Investigation Path Analysis, AI-powered Behavioral Analysis and Feedback

08

8.1 Operational Readiness as Evidence: Regulatory Alignment

From Documented Compliance to Demonstrable Readiness

A few operational indicators are sufficient to make this value tangible.

Breach cost and dwell time: According to IBM's 2025 *Cost of a Data Breach* report, organizations that keep the breach lifecycle under 200 days incur, on average, $1.1 million less in costs compared to longer-running incidents. The same report shows that human error is a contributing factor in 26% of breaches. These findings suggest that the quality of investigation under operational pressure plays a significant role in determining the overall impact of an incident.

Burnout and turnover cost: According to Sophos' 2025 research, cybersecurity professionals lose an average of 4.8 hours per week due to stress and burnout, representing an increase of more than 25% compared to 2024. Research by Tines on SOC analysts indicates that 71% of participants experience burnout, while 64% are considering changing jobs. Across the industry, the cost of replacing a technical professional is estimated at approximately 80% of their annual salary. Applied to a mid-level SOC analyst, this translates to a replacement cost in the range of $80,000 to $120,000 per departure. Additionally, ISC2's 2025 data indicates that it can take 4 to 9 months for a new analyst to become fully independent and productive.

Senior analyst capacity: In many SOC environments, senior analysts are required to spend a significant portion of their time mentoring and correcting the outputs of junior analysts. When investigation approaches and decision quality become more consistent, this burden is meaningfully reduced, allowing senior capacity to be redirected toward higher-priority threats.

Regulatory assurance cost: DORA has been in effect since January 2025, and NIS2 will significantly increase compliance pressure across the EU as of October 2026. Both frameworks require organizations not only to maintain policies and procedures, but also to demonstrate that these are effectively applied under real operational conditions. Under DORA, non-compliance penalties can reach up to 2% of global annual turnover. Having a validation mechanism that evidences readiness both accelerates audit preparation and strengthens assurance.

Taken together, these four factors form a clear picture: shorter breach lifecycles, reduced turnover, more efficient use of senior capacity, and lower compliance preparation costs. Combined, they establish a strong and defensible business case for GRASP—one that does not rely on speculative performance claims, but on observable trends supported by industry data.

Compliance and assurance frameworks have traditionally focused on documented processes, assigned responsibilities, and formal control structures. These remain necessary. However, expectations are shifting from the mere existence of documentation to demonstrable operational capability. A readiness validation approach such as GRASP does not replace governance, procedures, or compliance artifacts; it strengthens them with evidence generated at the point where policy meets execution. When an organization can demonstrate not only that it has an incident response plan, but also that it observes and improves human performance under pressure, it achieves a stronger assurance posture.

Compliance and assurance frameworks have traditionally focused on documented processes, assigned responsibilities, and formal control structures. These remain necessary. However, expectations are shifting from the mere existence of documentation to demonstrable operational capability. A readiness validation approach such as GRASP does not replace governance, procedures, or compliance artifacts; it strengthens them with evidence generated at the point where policy meets execution. When an organization can demonstrate not only that it has an incident response plan, but also that it observes and improves human performance under pressure, it achieves a stronger assurance posture. As expectations continue to shift toward demonstrable operational capability, the ability to observe and develop human performance under realistic conditions becomes not just a learning advantage, but a fundamental component of how organizational readiness is justified.

8.1 Operational Readiness as Evidence: Regulatory Alignment

Framework	Requirement	What is Expected	GRASP Contribution
DORA	Article 11 – ICT Incident Management	Effective detection and response under real conditions	Makes human performance observable under pressure
DORA	Article 13 – Testing	Regular resilience testing	Enables scenario-driven readiness validation
DORA	Article 14 – Post-Incident Review	Continuous improvement	Turns behavior into feedback signals
DORA	Articles 26–27 – TLPT	Threat-led testing including people	Extends testing to human decision-making
NIS2	Article 21(2)(f) – Effectiveness	Measure effectiveness of controls	Provides measurable readiness signals
NIS2	Article 23 – Reporting	Timely and accurate reporting	Improves escalation and clarity
PCI DSS 4.0	Req. 12.10 – Incident Response	Tested IR capability	Validates real-world performance
ISO 27001	A.5.24 – Incident Management	Consistent IR processes	Makes performance measurable
NIST CSF	Respond / Improve	Continuous improvement	Enables feedback-driven development

Turn every incident into measurable improvement

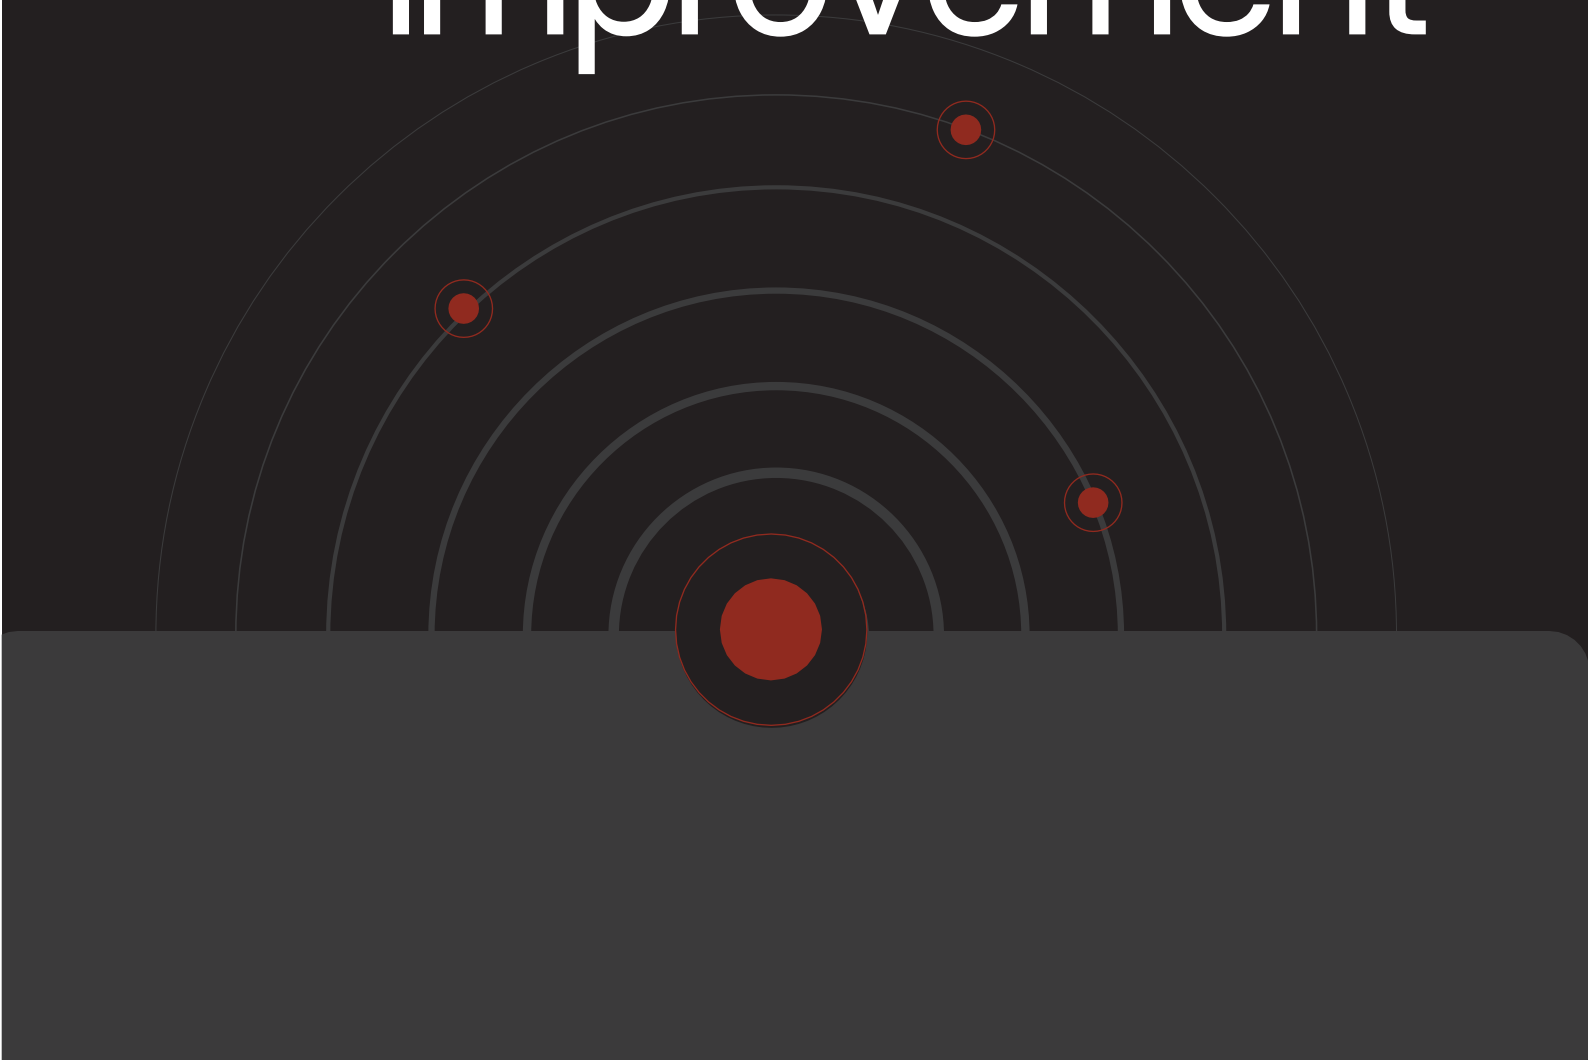

Appendix

09

Human Readiness: The ability of analysts and response teams to perform effectively in real incidents, particularly under uncertainty, time pressure, and operational stress.

Functional Readiness: The behavioral manifestation of readiness: context formation, hypothesis generation, model updating, prioritization, escalation, and decision quality under pressure.

Decision Quality: The accuracy, timing, and contextual appropriateness of decisions made under uncertainty. The most direct output of incident response performance.

Hypothesis Formation: The process of constructing meaningful, testable, and guiding working hypotheses from incomplete and fragmented signals.

Prediction Error: The mismatch between existing expectations or hypotheses and new evidence. The core driver of learning and adaptation.

Adaptive Response: The ability to update, refine, or shift the working model in response to conflicting evidence and changing context.

Context Formation: The process of transforming incomplete and fragmented signals into a functional model of what is happening.

Investigation Flow: The structured sequence of hypotheses, decisions, and actions followed by an analyst throughout an incident.

Behavioral Visibility: The ability to observe not only the outcome, but the full investigation process leading to it.

Escalation Quality: The accuracy and timing of decisions regarding when and how to escalate an incident.

Operational Pressure: The combined impact of time pressure, interruptions, communication load, uncertainty, and stakeholder attention on response performance.

Core Concepts

Term	Description
Time to First Working Hypothesis	Time required for the analyst to form the first actionable and guiding hypothesis
Hypothesis Stability	The extent to which the initial hypothesis is maintained or frequently revised
Investigation Coherence	The logical consistency and progression of the investigation over time
Pivot Quality	The accuracy and contextual relevance of directional shifts during investigation
Adaptation to Contradictory Evidence	The speed and quality of updating hypotheses when faced with conflicting evidence
Escalation Timing	The correctness and timing of escalation decisions
Prioritization Under Ambiguity	The ability to prioritize meaningful actions under uncertainty
Signal-to-Noise Ratio	The proportion of analyst output that contributes to meaningful progress
Unnecessary Branching	The degree of deviation into irrelevant or low-value investigation paths
Communication Under Pressure	The clarity, accuracy, and effectiveness of communication under stress

Example Readiness Signals

References

- Clark, A. (2013). *Whatever next? Predictive brains, situated agents, and the future of cognitive science*. Behavioral and Brain Sciences, 36(3), 181–204.
- Federal Bureau of Investigation (FBI). (2025). *2024 Internet Crime Report*. Internet Crime Complaint Center (IC3).
- Friston, K. (2010). *The free-energy principle: A unified brain theory?* Nature Reviews Neuroscience, 11(2), 127–138.
- IBM Security & Ponemon Institute. (2025). *Cost of a Data Breach Report 2025*.
- ISC2. (2025). *Cybersecurity Workforce Study 2025*.
- Kolb, D. A. (1984). *Experiential Learning: Experience as the Source of Learning and Development*. Prentice Hall.
- Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.
- SANS Institute & GIAC. (2025). *Cybersecurity Workforce Research Report*.
- Sophos. (2025). *The Human Cost of Vigilance: Addressing Cybersecurity Burnout in 2025*.
- Tines. (2023). *Voice of the SOC Report*.
- Cynet. (2023). *CISO Stress Survey*.
- High5Test. (2025). *Employee Turnover Statistics (2024–2025)*.
- MITRE. (2024). *ATT&CK Framework*.
- ENISA. (2024). *NIS2 Implementation Guidance*.
- Cyber Struggle. (n.d.). *Ranger Program Observational Dataset (23 cohorts, 20.000+ Exercises)*. Internal dataset based on longitudinal simulation observations; collected without a control group and does not assert causal inference.

Cyber Struggle

TRAINING THE SPECIAL FORCES OF CYBER WORLD

S46 GRASP
Adversarial Readiness of
Incident Response and SOC Teams