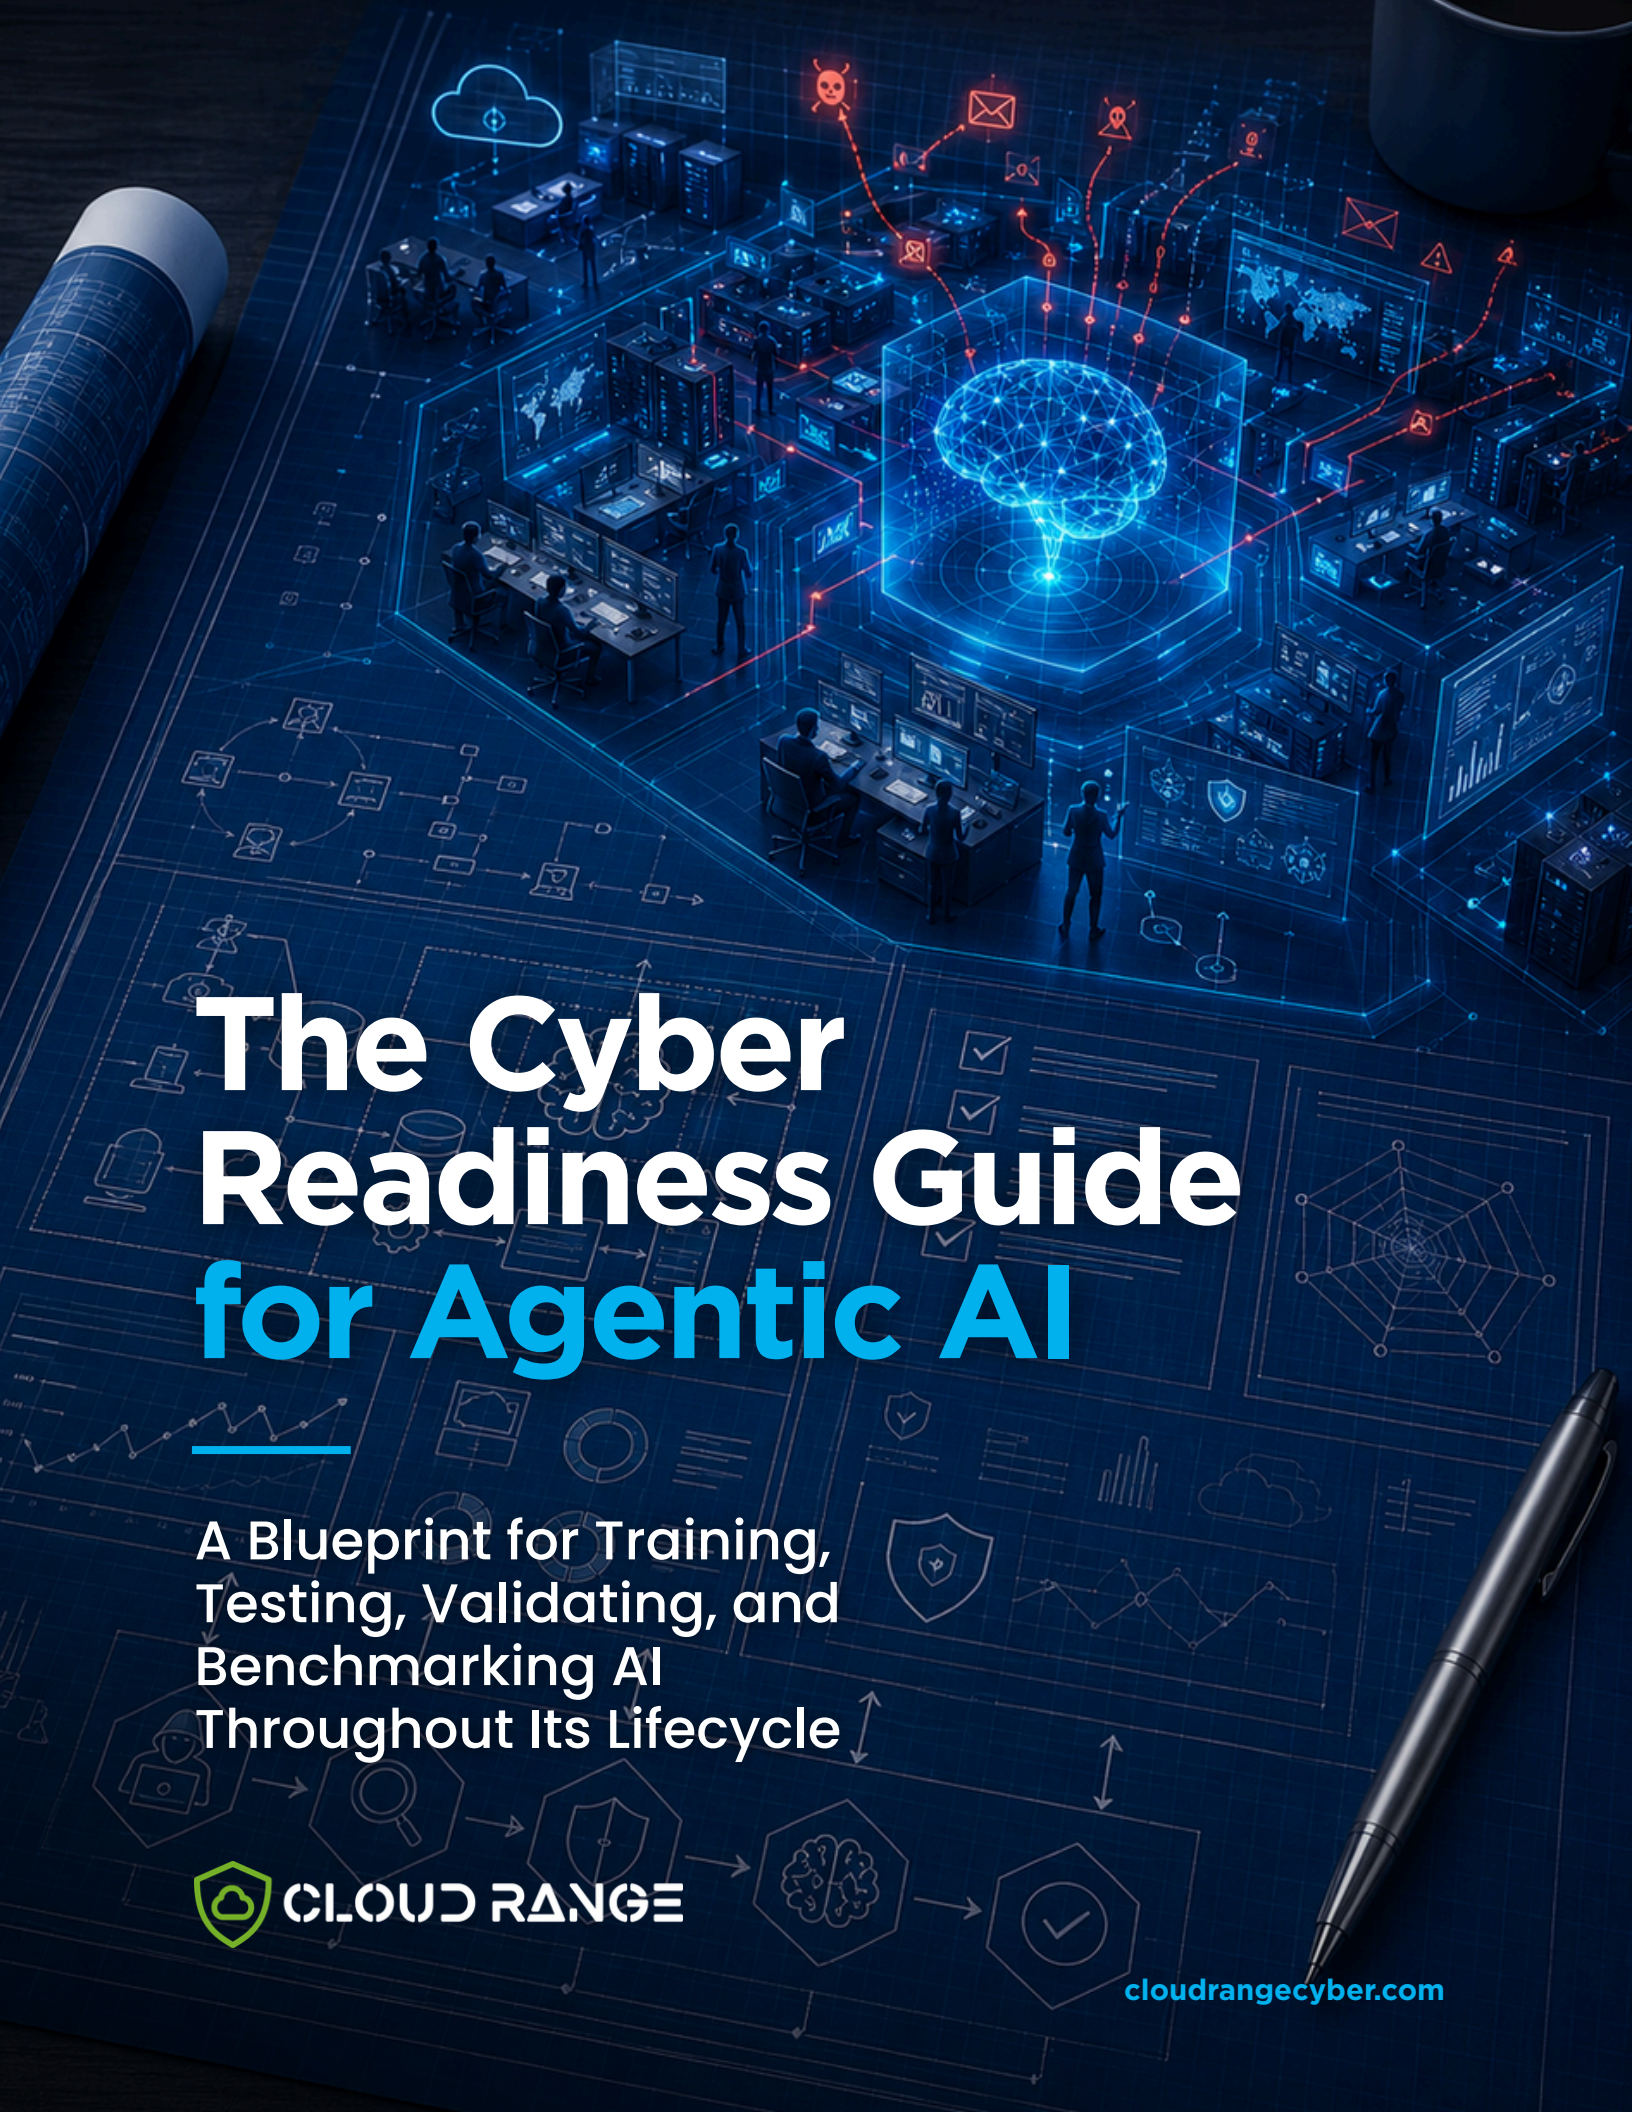

The Cyber Readiness Guide for Agentic AI

A Blueprint for Training, Testing, Validating, and Benchmarking AI Throughout Its Lifecycle

CLOUD RANGE

cloudrangecyber.com

Executive Summary

Artificial intelligence is rapidly becoming part of modern security operations. Organizations are integrating AI into alert triage, investigations, enrichment workflows, response processes, and offensive security activities. At the same time, vendors are introducing increasingly autonomous agents capable of performing tasks that were previously reserved for human analysts.

AI promises greater efficiency, scale, and speed across security operations. As pressure mounts to do more with limited resources, many security leaders view AI as a critical component of the future SOC.

However, AI adoption is accelerating while meaningful validation remains rare.

Many organizations deploy AI after functionality testing, proof-of-concept projects, or limited pilot programs without a clear understanding of how the system will perform when operational risk is on the line. Functionality may be proven, but readiness often remains unverified.

This creates a growing readiness gap. As AI gains access to enterprise data, security tools, workflows, and decision-making responsibilities, security leaders need a more rigorous way to evaluate performance, understand limitations, and determine where human oversight remains necessary.

This ebook presents a framework for agentic SOC validation. It explores the risks of learning in production, outlines the components of effective AI validation, and provides security leaders with a practical approach for evaluating readiness before expanding AI responsibility.

Because AI deployment is not simply a technology decision.

It is a risk management decision.

Table of Contents

INTRODUCTION

AI Is Being Trusted Before It's Proven

The cybersecurity industry is entering a period of rapid transformation. AI is moving from experimental projects into operational workflows, and organizations are increasingly relying on AI systems to support investigations, prioritize alerts, analyze telemetry, and assist with response activities. In some cases, AI agents are beginning to perform tasks with minimal human involvement, creating new opportunities for efficiency while introducing new forms of operational risk.

The challenge is not whether AI can provide value. The challenge is understanding how it behaves when conditions become complex, uncertain, or adversarial.

Many organizations can explain what their AI systems are designed to do. Far fewer can explain how those systems will perform when presented with incomplete information, conflicting data, unfamiliar attack techniques, or unexpected operating conditions. Even fewer have a structured process for evaluating those behaviors before AI is trusted with meaningful responsibility.

This gap between functionality and readiness represents one of the most significant challenges facing security leaders today. Traditional testing environments often confirm that a model or agent can complete a task. They rarely reveal how the system behaves when conditions become noisy, dynamic, or adversarial. As a result, organizations frequently discover limitations, weaknesses, and failure modes after deployment rather than before it.

That approach may be acceptable when evaluating a new productivity tool. It becomes significantly more risky when AI is integrated into security operations and granted access to sensitive systems, enterprise data, operational workflows, and security decision-making processes.

Security leaders remain accountable for the outcomes of those systems. Whether an action is taken by a human analyst or an autonomous agent, responsibility for managing risk does not change. Organizations that cannot explain how AI behaves, where it fails, or when oversight remains necessary are accepting risk they do not fully understand.

Understanding AI behavior before it is trusted—and continuously validating that behavior as threats, environments, and workflows evolve—is becoming an essential component of modern cyber readiness.

AI Doesn't Need More Trust.

It Needs More Validation.

AI adoption is accelerating across security operations. The question is no longer whether organizations will use AI. The question is whether they understand how it will perform when operational risk is on the line.

Trust is assumed.
Validation is proven.

CHAPTER 1

The Danger of Learning in Production

Organizations have established processes for validating people, technologies, and workflows before trusting them in operational environments.

AI should be no different.

Yet many organizations are deploying AI into security operations while simultaneously learning how it behaves. Functionality is tested, pilot programs are launched, and AI begins interacting with real systems, real data, and real workflows before leaders fully understand its limitations, failure modes, or operational risks. This creates a readiness gap that many organizations do not discover until AI is already influencing security outcomes.

An AI system may perform well during demonstrations, proof-of-concept projects, or limited testing. That does not necessarily mean it is prepared for the complexity and unpredictability of security operations. Conflicting telemetry, incomplete information, evolving attack techniques, unexpected user behavior, and operational noise often expose risks that were never apparent during initial evaluations. As a result, organizations frequently discover how their AI behaves only after it has been trusted with meaningful responsibility.

WHAT HAPPENS WHEN YOU LEARN IN PRODUCTION?

Organizations often deploy AI into production and then discover what it can access, how it behaves, and where it fails.

At that point, they are no longer validating AI. They are learning in production.

The cost of those lessons depends on what the AI has been trusted to do and how quickly failures can be detected and corrected.

Unlike traditional software, AI systems do not always respond predictably when conditions change. They make decisions based on patterns, context, probabilities, and inference rather than fixed rules. That flexibility is one of AI's greatest strengths, but it also makes behavior more difficult to predict as environments become more complex.

An AI agent may perform exceptionally well during testing yet struggle when confronted with unfamiliar data, contradictory signals, or circumstances that differ from its training environment. The issue is not whether AI will fail. The issue is whether organizations understand its failure modes before those failures affect production operations.

LEARNING THE WRONG LESSONS

Consider a common scenario: An organization deploys an AI-enabled workflow to help analysts investigate alerts and gather supporting information. Initial testing is successful. Investigations are completed faster, analysts receive useful context, and productivity improves.

Weeks later, an analyst submits what appears to be a routine request. In responding, the AI retrieves information from a location it was never expected to access. Sensitive operational information, administrative documentation, or privileged credentials are exposed—not because the AI was compromised, but because nobody anticipated how it would behave under that specific set of conditions.

The organization learns something important about the AI. Unfortunately, it learns it in production.

This is one of the most significant challenges facing security leaders today. AI systems are frequently evaluated based on whether they work rather than whether they are ready. When readiness is assumed, organizations often discover limitations only after those limitations have operational consequences.

These failure modes are not theoretical. They align closely with risks identified in the OWASP Top 10 for LLM Applications, including excessive agency, prompt injection, and sensitive data disclosure. The challenge is not whether these risks exist. The challenge is understanding how an AI system responds before those failures affect production operations.

AGENTIC AI RAISES THE STAKES

The risks become even greater as organizations move beyond AI assistants and toward agentic AI.

Unlike systems that simply generate recommendations, agentic AI can be granted authority to interact with security tools, access enterprise data, perform investigations, initiate workflows, and support response activities. As organizations pursue greater efficiency and autonomy, AI systems are increasingly being trusted with responsibilities that were once reserved for human analysts.

This shift fundamentally changes the risk equation. As AI gains responsibility, organizations need more than confidence that it works. They need evidence that it will behave as expected when conditions change.

A recommendation that is wrong can be reviewed and corrected. An automated action that is wrong may have a much greater impact. The more responsibility organizations assign to AI, the more important it becomes to understand how that system performs, where its limitations exist, and what level of oversight remains necessary.

VALIDATION CHANGES THE CONVERSATION

The goal of validation is not simply to prove that AI succeeds. It is to develop a clear understanding of strengths, limitations, failure modes, and readiness so organizations can make informed decisions about deployment, oversight, and risk.

Security leaders should be able to answer questions such as:

- Under what conditions does the AI perform well?
- Under what conditions does performance degrade?
- What actions require human approval?
- What level of autonomy is appropriate?
- What risks are acceptable?
- What risks are not?

These are not technology questions. They are risk management questions.

Organizations that cannot answer them are not deploying AI with confidence. They are deploying AI with assumptions.

And assumptions are a poor foundation for security operations.

CHAPTER 2

What Security Leaders Must Understand About AI Risk

AI is changing security operations in ways that extend far beyond automation. Organizations are increasingly using AI to support investigations, analyze telemetry, prioritize alerts, enrich data, recommend actions, and perform tasks that previously required human judgment. As AI becomes more deeply integrated into security workflows, it is also gaining access to enterprise systems, sensitive information, and operational decision-making processes.

The opportunity is significant. So is the risk.

Many organizations approach AI primarily as a technology initiative. Security leaders should view it as a risk management challenge. The question is not simply whether AI can perform a task. The question is what level of responsibility it has been given, what actions it can influence, and what happens when it gets something wrong.

AI IS BEING TRUSTED WITH MORE RESPONSIBILITY

The evolution of AI in security operations is not defined by what AI knows. It is defined by what AI is trusted to do.

Early AI systems typically generated recommendations that analysts could review before taking action. Today's agentic AI systems are increasingly being trusted to perform investigations, interact with security tools, execute workflows, and support response activities with varying levels of human involvement.

As organizations pursue greater efficiency and automation, AI is moving closer to operational decision-making. The more responsibility AI receives, the more important it becomes to

understand how it behaves, where its limitations exist, and what level of oversight remains appropriate.

AI DECISIONS MUST BE OBSERVABLE

Accountability is essential to effective security operations. When analysts make decisions, organizations can typically review the evidence, understand the reasoning, and determine why a particular action was taken. AI introduces a new challenge because many organizations can observe what an AI system did but cannot fully explain why it did it.

This lack of visibility creates risk. Security leaders need to understand what information influenced a decision, how recommendations were generated, and whether the AI is operating in a way that aligns with organizational objectives. Without that visibility, it becomes difficult to identify weaknesses, correct mistakes, or understand how performance may change as conditions evolve. Validation helps close that gap by providing insight into behavior, not just outcomes.

AI IS A NON-HUMAN IDENTITY

Every new AI agent effectively becomes another identity operating within the environment.

Like human users, AI agents may have access to systems, applications, data, workflows, and security tools. Unlike human users, they can often operate at machine speed, interact with multiple systems simultaneously, and perform actions at a scale that is difficult to monitor through traditional approaches.

This creates new governance and security challenges. Organizations must understand what AI agents can access, what permissions they have been granted, what actions they can perform, and how those activities are monitored. As the number of AI agents grows, security leaders must treat them as operational identities that require the same level of visibility, oversight, and validation as any other privileged entity in the environment.

> AI risk is created by access, authority, and autonomy.
>
> That's why it's critical to understand how it behaves, where it fails, and what level of oversight remains necessary.

MORE RESPONSIBILITY CREATES MORE RISK

The risks associated with AI are often discussed in terms of autonomy, but autonomy is only part of the equation.

Risk is created whenever AI is given access to sensitive information, authority to influence decisions, or the ability to perform actions that affect security outcomes. Even systems that do not operate autonomously can create significant risk if they are trusted to provide inaccurate recommendations, expose sensitive data, or influence critical decisions.

As organizations expand AI responsibilities, they must also expand their understanding of the risks associated with those responsibilities. The goal is not to eliminate risk. The goal is to understand it well enough to determine where AI creates value, where additional safeguards are required, and where human oversight remains necessary.

EVERY AI HAS FAILURE MODES

Every security technology has limitations, and AI is no different. The challenge is that many organizations do not fully understand those limitations until after deployment. AI may produce inaccurate conclusions, misinterpret instructions, rely on flawed assumptions, expose information it should not access, or behave unpredictably when confronted with unfamiliar situations. These outcomes are not necessarily indicators of poor technology. They are realities that must be understood and managed.

The question is not whether AI will fail. The question is whether organizations understand its failure modes before those failures affect production operations.

Organizations that understand failure modes can manage risk. Organizations that discover them during an active incident are managing consequences.

CLOUD RANGE

AI VALIDATION READINESS CHECKLIST

10 Questions Every Security Leader Should Be Able to Answer

Use this checklist to evaluate your internal engineering teams and external AI vendors before deployment. Before expanding the role of AI within security operations, leaders should be able to answer the following questions with confidence.

1. What data can the AI access?

Does the AI have access only to the information required for its function, or can it reach additional systems, repositories, applications, and data sources?

2. What tools can the AI interact with?

Can it view information only, or can it modify configurations, execute workflows, trigger actions, or interact with security controls?

3. What permissions has the AI been granted?

Have least-privilege principles been applied, or does the AI have broader access than necessary?

4. How does the AI behave when it receives incomplete, incorrect, or conflicting information?

Has the organization evaluated how the system responds when the data it receives is imperfect, ambiguous, or misleading?

5. What are the AI's known failure modes?

Can the organization identify situations where the AI struggles, performs inconsistently, or produces undesirable outcomes?

6. How will AI activity be monitored?

Can security teams clearly see what actions the AI is taking, what decisions it is making, and what resources it is accessing?

7. What requires human approval?

Has the organization clearly defined where human-in-the-loop oversight remains mandatory?

8. How will performance be measured?

Can the organization determine whether the AI is achieving desired outcomes, operating consistently, and performing as expected over time?

9. Where are the boundaries of acceptable risk?

Has the organization defined clear limits around autonomy, authority, access, and operational decision-making?

10. How will validation continue over time?

What process exists to evaluate performance as threats, environments, workflows, and AI systems evolve?

If these questions cannot be answered, the organization may be trusting AI before it has been properly validated.

CHAPTER 3

Validation Is More Than Testing

Most organizations test AI before deployment. *Fewer validate it.*

The difference is significant because testing and validation answer fundamentally different questions. Testing is designed to determine whether an AI model, agent, or workflow functions as intended. It confirms that the system can complete specific tasks, generate expected outputs, and perform within defined parameters. While testing is necessary, it provides only a partial view of readiness.

Validation goes further. It evaluates how AI performs when confronted with uncertainty, changing conditions, conflicting information, and adversarial activity. Rather than asking whether the AI works, validation asks whether the AI can be trusted to perform consistently and predictably when operational risk is on the line.

As AI moves deeper into security operations, security leaders are not simply deploying technology; they are assigning responsibility. Before that responsibility expands, organizations need evidence that AI can perform reliably across a range of conditions, not just in controlled demonstrations or ideal scenarios.

VALIDATION HAPPENS UNDER PRESSURE

Security operations rarely occur under perfect conditions. Analysts routinely make decisions with incomplete information, conflicting signals, changing priorities, and active threats. AI systems must operate in those same environments. As a result, effective validation requires

more than predefined test cases and expected outcomes. It requires exposing AI to the kinds of situations it will encounter when security outcomes are at stake.

This is where many organizations encounter a readiness gap. Traditional testing often evaluates AI using known inputs, predictable workflows, and controlled environments. Those exercises can confirm functionality, but they rarely reveal how performance changes when complexity increases, assumptions break down, or unexpected conditions emerge.

Effective validation intentionally introduces pressure. AI should be evaluated against different attack techniques, varying levels of operational complexity, changing workflows, incomplete data, and unexpected conditions. Organizations should understand how the AI responds when information becomes unreliable, priorities shift, or adversaries behave in ways the system has not previously encountered.

Furthermore, true validation requires context. An AI agent cannot reliably identify anomalous behavior until it understands what normal network traffic looks like within your specific environment. True validation requires the ability to ingest your organization's own packet capture (PCAP) files or generated traffic baselines into the testing environment, rather than relying on generic, out-of-the-box lab data.

The objective is not to create failure for the sake of failure. The objective is to understand performance boundaries before those boundaries are discovered in production.

Testing confirms functionality. Validation proves readiness.

TESTING	VS.	VALIDATION
Point-in-time activity		Ongoing process
Expected outcomes		Expected and unexpected outcomes
Focuses on features		Focuses on operational performance
“Did it work?”		"Can we trust it?"
Activity		Evidence

VALIDATION REVEALS STRENGTHS AND LIMITATIONS

One of the most common misconceptions about validation is that it exists primarily to identify problems. In reality, validation provides a much more complete picture of performance.

Organizations need to understand where AI performs well, where it creates value, and which tasks it can execute consistently. They also need to understand where performance begins to degrade, what conditions create risk, and where additional oversight may be required. Both perspectives are essential for making informed decisions about deployment, responsibility, and operational use.

Validation also helps establish what successful performance should look like. Neither speed nor accuracy alone determines readiness. In many cases, the most valuable insight comes from understanding the relationship between speed, accuracy, consistency, decision quality, and risk. These factors help organizations define meaningful performance expectations rather than relying on assumptions, isolated successes, or vendor claims.

The result is a clearer understanding of how AI performs, where limitations exist, and how the technology should be used within security operations.

> The goal is not to prove the AI works.
>
> The goal is to understand **how** it works, **where** it performs well, where it **doesn't**, and what **success** actually looks like.

VALIDATION CREATES EVIDENCE

Security leaders are ultimately responsible for decisions involving risk, and those decisions become more difficult when they are based on assumptions rather than evidence. Validation provides the data needed to make informed decisions about AI deployment, oversight, and operational responsibility. It helps organizations understand whether performance aligns with expectations, whether identified risks are acceptable, and whether the AI is prepared for greater responsibility.

Most importantly, validation replaces uncertainty with measurable evidence. Organizations gain a clearer understanding of how AI behaves, how it performs, and where confidence is justified. That evidence becomes the foundation for readiness and allows leaders to make decisions based on performance rather than assumptions.

Readiness is not established through testing alone. Readiness is established through validation.

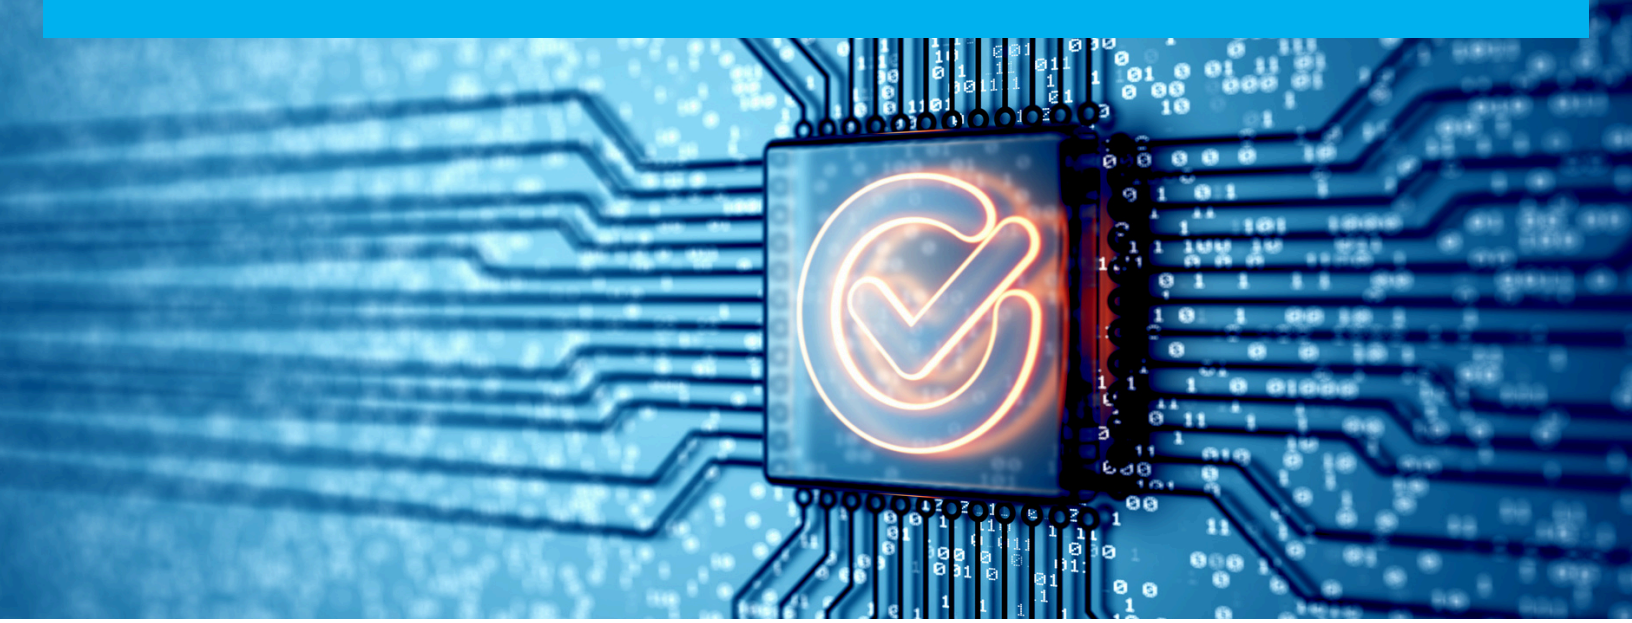

CHAPTER 4

How to Prove AI Readiness

Security leaders are asked to make decisions about deployment, autonomy, oversight, and risk. Those decisions become difficult when they are based on assumptions rather than evidence.

The challenge is not determining whether AI can perform a task. The challenge is determining whether it can perform consistently, reliably, and within acceptable risk boundaries. Before organizations expand AI responsibilities, they need confidence that performance aligns with expectations and that limitations are understood.

Readiness goes beyond functionality. The question is not whether AI can complete a task, but whether it can be trusted with responsibility.

READINESS REQUIRES EVIDENCE

Many organizations evaluate AI based on isolated demonstrations, successful pilot programs, or limited testing exercises. While those activities may provide useful information, they rarely provide enough evidence to support decisions about operational responsibility.

Readiness requires a broader view of performance. Security leaders need to understand how AI behaves across different scenarios, how consistently it performs over time, and where limitations begin to emerge. They also need visibility into the tradeoffs between speed, accuracy, decision quality, and risk.

Without that evidence, deployment decisions are often based on assumptions, vendor claims, or isolated successes. Validation replaces those assumptions with measurable data that can be used to support governance, oversight, and risk management decisions.

WHAT VALIDATION SHOULD REVEAL

Validation goes beyond functionality. It reveals how AI is likely to perform under real-world conditions—giving organizations the evidence needed to deploy with confidence.

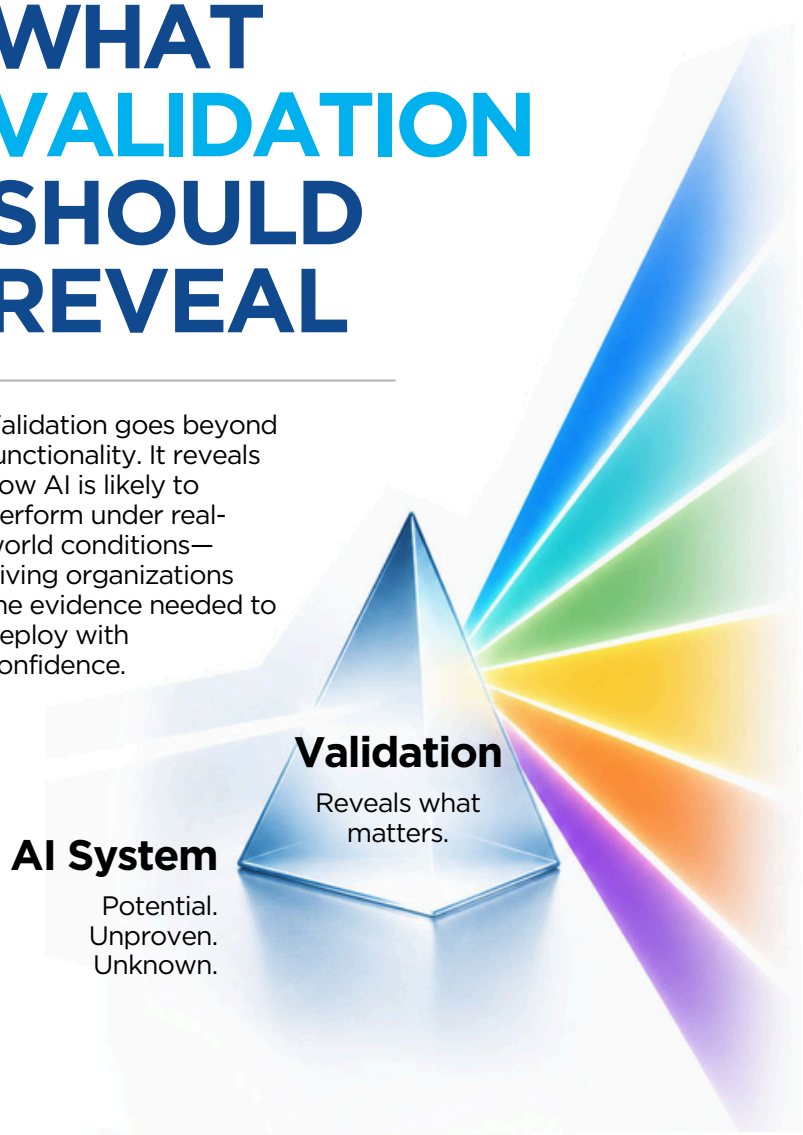

Accuracy

Measures the correctness and reliability of AI decisions across a range of realistic scenarios.

Performance

Evaluates how effectively the AI detects, responds, and completes tasks under operational conditions.

Consistency

Determines whether the AI delivers reliable results over time and with different inputs and environments.

Limitations

Identifies failure modes, blind spots, and situations where human oversight remains necessary.

Risk

Evaluates the potential operational, security, and business impact of AI failures or misuse.

Readiness

Determines whether the AI is prepared to operate safely and effectively in its intended role.

DEFINING SUCCESS

One of the most valuable outcomes of validation is a clear understanding of what success actually looks like.

Organizations often assume they will recognize success when they see it. In practice, meaningful performance expectations must be defined and measured. Speed alone does not determine readiness. Nor does accuracy alone. An AI system that responds quickly but produces unreliable recommendations may create more risk than value. Likewise, an AI that achieves high accuracy but cannot operate effectively at scale may not meet operational requirements.

Validation helps organizations establish performance expectations that align with business objectives, operational realities, and risk tolerance. Those expectations become the basis for evaluating readiness and making informed decisions about how AI should be used.

HUMAN AND AI PERFORMANCE SHOULD BE MEASURED TOGETHER

Understanding AI performance in isolation provides only part of the picture. One of the most effective ways to evaluate readiness is to compare AI performance against human performance in the same scenarios.

Human versus AI benchmarking provides visibility into where AI creates value, where human expertise remains essential, and how both can work together most effectively. In some situations, AI may outperform human analysts in speed, scale, or consistency. In others, human judgment, context, and decision-making may remain critical to achieving desired outcomes.

The objective is not to determine whether AI or humans are better. The objective is to understand the strengths and limitations of each and identify the most effective balance between automation and human oversight.

These comparisons also help organizations determine where human-in-the-loop review remains necessary and where greater autonomy may be appropriate.

READINESS IS A RISK DECISION

Security leaders are responsible for determining how much authority AI should receive, what safeguards should remain in place, and what level of risk is acceptable. Those decisions require more than confidence. They require evidence.

Validation provides that evidence. It enables organizations to make informed decisions about deployment, oversight, responsibility, and risk based on measurable performance rather than assumptions.

Ultimately, readiness is not defined by what the AI is capable of doing. Readiness is defined by what the organization can confidently trust it to do.

CHAPTER 5

Validation Never Ends

Many organizations approach AI validation as a milestone. A model is evaluated, a pilot is completed, performance is reviewed, and a deployment decision is made. At that point, validation is often considered complete.

The problem is that AI systems do not operate in static environments. Models are updated, data sources change, workflows evolve, business priorities shift, and threat actors continuously adapt their tactics. The conditions that influence performance today may look very different six months from now.

AI systems themselves may also change over time. Models are retrained, agents learn from new information, integrations are added, permissions expand, and workflows evolve. These changes can improve performance, but they can also introduce new risks, create unintended behaviors, or expose limitations that were not present during earlier evaluations.

An AI system that performs well today may perform differently tomorrow. New attack techniques may expose weaknesses that were never identified during previous evaluations. Changes to permissions, integrations, workflows, or operating environments can alter behavior in ways that affect performance, accuracy, and risk.

This is why validation should be viewed as an operational discipline rather than a deployment activity.

VALIDATION MUST KEEP PACE WITH CHANGE

AI exists in a constantly changing environment. New threats emerge. New attack techniques appear. Business processes evolve. New tools, data sources, integrations, and AI capabilities can all influence how systems behave and perform.

Organizations should continuously challenge assumptions, re-evaluate performance, and gather evidence that reflects current operating conditions rather than historical results. The goal is not to repeatedly prove the same outcome. The goal is to confirm that performance remains aligned with expectations as circumstances change.

VALIDATION SUPPORTS CONFIDENCE

Trust is often discussed as though it is a permanent state. In practice, confidence must be earned and maintained.

Security leaders need current, defensible evidence about how AI performs, where limitations exist, and what level of oversight remains necessary—not just for internal teams, but for board-level reporting and regulatory compliance.
Continuous validation helps organizations replace uncertainty with measurable performance data, allowing deployment, autonomy, and governance decisions to evolve alongside the technology itself.

Ultimately, AI changes how security work gets done, but it does not eliminate accountability. Security leaders remain responsible for the outcomes of the systems they deploy, the risks they accept, and the decisions they make.

Validation provides the evidence needed to make those decisions with confidence—not once, but continuously.

> The organizations with the most confidence in their AI are the ones that validate it continuously

Putting Continuous Validation Into Practice

Understanding why validation matters is only the first step. The bigger challenge is operationalizing it.

Cloud Range's AI Validation Range™ helps organizations move beyond functionality testing and pilot programs by providing a controlled enterprise environment to train, test, validate, benchmark, and continuously measure AI performance before and after deployment.

Built on the broader Cloud Range platform, AI Validation Range combines realistic cyber ranges, live-fire simulations, AI validation, objective performance measurement, and continuous readiness into a single managed environment.

The platform replicates realistic enterprise infrastructure, security tools, users, workflows, attack activity, network traffic, and operational complexity without introducing risk into production.

Organizations can also incorporate packet capture (PCAP) files and generated traffic baselines so AI learns what "normal" looks like within their own environment before being challenged with realistic attacks. This provides far more meaningful validation than generic lab environments or static datasets.

Organizations use AI Validation Range to:

- Validate AI models before and after deployment
- Train and evaluate defensive SOC agents
- Develop and assess offensive security agents
- Test AI behavior against live-fire cyberattacks
- Evaluate accuracy, consistency, and decision quality
- Identify failure modes and operational limitations
- Benchmark human and AI performance side by side
- Measure readiness using repeatable scenarios and performance data

Because the environment is fully controlled, organizations can safely challenge AI systems with different attack techniques, operational workflows, data conditions, and adversarial activity to understand how performance changes as complexity increases.

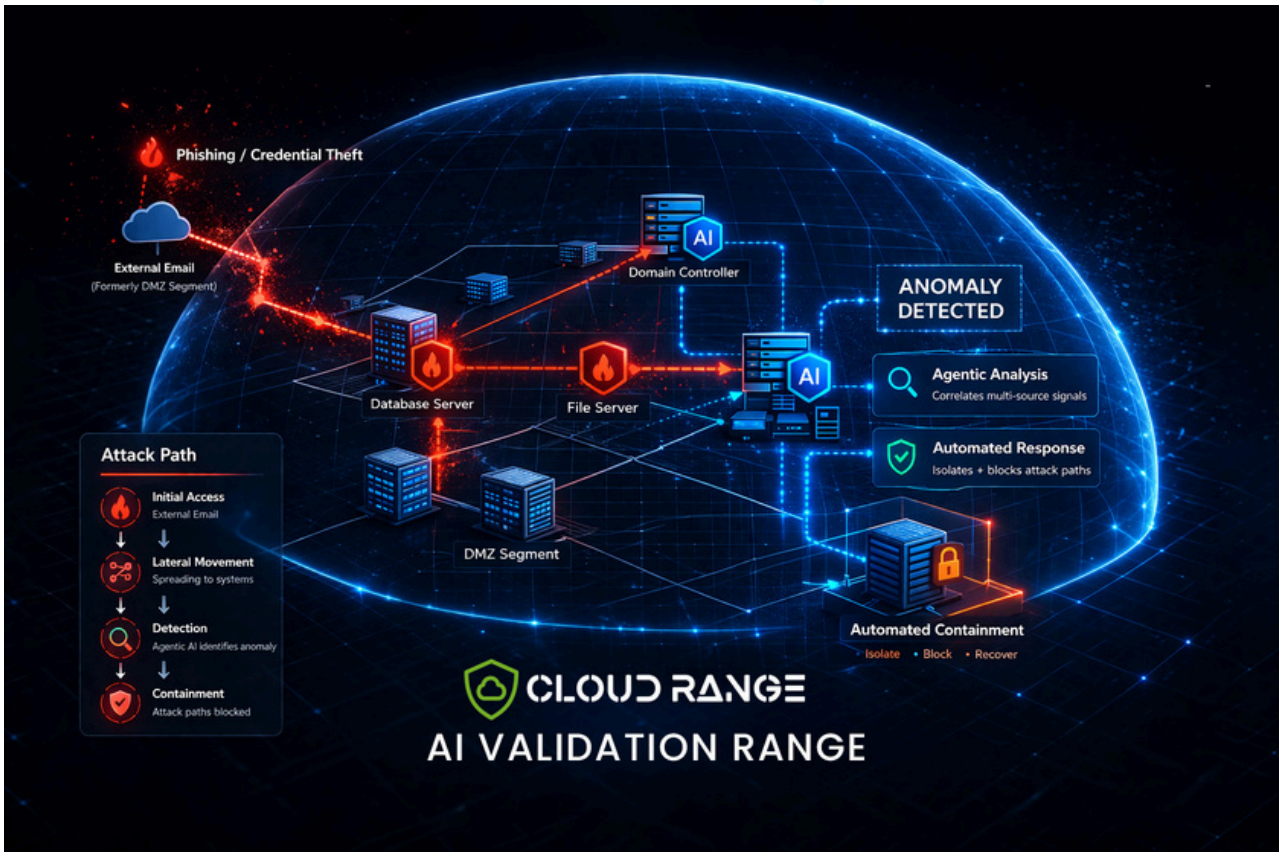

VALIDATE AI MODELS AND AGENTS

Evaluate commercial, open-source, and proprietary AI for accuracy, reliability, prompt manipulation, sensitive data exposure, and operational behavior. Validation reveals how models perform under realistic conditions, where limitations exist, and whether they are ready for deployment.

TRAIN AND VALIDATE AGENTIC AI

Measure defensive and offensive AI agents against live-fire cyberattacks aligned to real-world adversary behaviors and frameworks such as MITRE ATT&CK®. Organizations gain objective evidence of how agents investigate, respond, adapt, and perform under pressure.

BENCHMARK HUMAN AND AI PERFORMANCE

Compare analysts and AI against identical scenarios to understand where automation accelerates response, where human judgment remains essential, and how both work together most effectively.

SUPPORT CONTINUOUS VALIDATION

Validation shouldn't end after deployment. Continuously measure readiness as AI systems, threats, workflows, and environments evolve to ensure performance remains aligned with expectations over time.

Organizations should not have to guess how their AI will perform when it matters. Cloud Range helps them prove it.

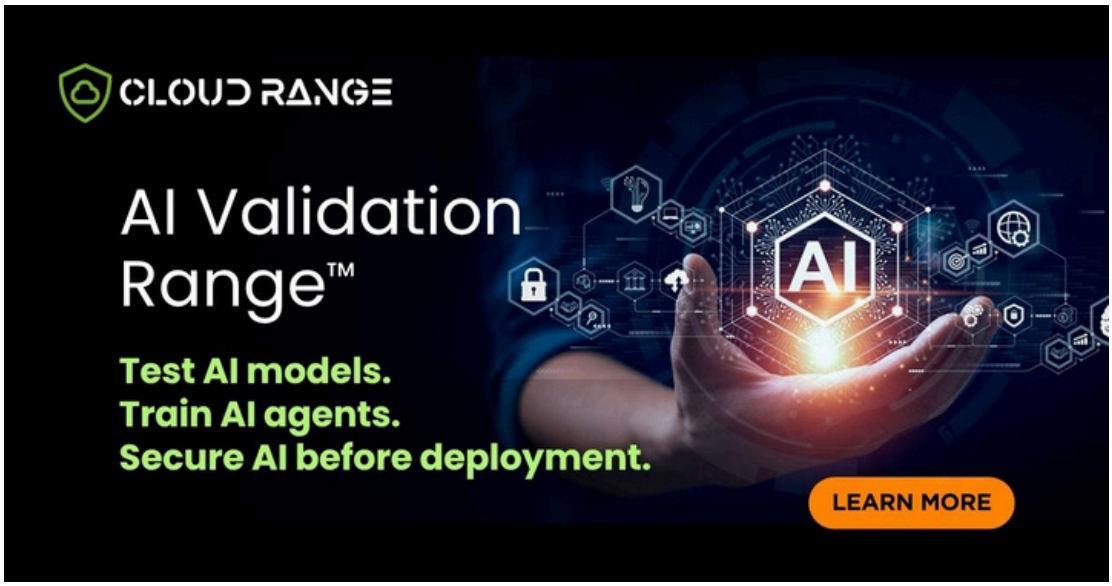

ABOUT CLOUD RANGE

Cloud Range's revolutionary cloud-based cyber range platform helps organizations measurably reduce cyber risk by testing and training both human and AI SOC agents to defend against real attacks. Its simulation platform includes an ever-changing library of realistic, live-fire attacks, IT/OT/IoT/cloud environments, skill development labs, advanced tabletop exercises, reporting, and more. Its AI Validation Range enables organizations to safely test, train, and validate AI models and agents without exposing real data in production environments. Ensure your organization is battle-ready.

CLOUD RANGE

START PROVING AI READINESS.
REQUEST A DEMO.

Contact Cloud Range at
info@cloudrangecyber.com

cloudrangecyber.com