



Agentic AI Assurance for FedRAMP 20x

A nine-indicator augmentation framework aligned to the Blueprint Alliance architecture and grounded in NIST, OWASP, and OMB guidance

38North Security · October 2026

Author: Austen Summers

1. Executive Summary

FedRAMP 20x uses Key Security Indicators (KSIs) to evaluate the security posture of cloud services. It does not yet define agent-specific indicators for software that can plan, call tools, delegate work, and act on systems with limited human direction.

Agentic systems introduce assurance questions that traditional control evidence does not fully answer. What authority does an agent actually have? What actions did it take? Can it delegate access or instructions to another agent? Can consequential effects be stopped before they occur? And when something goes wrong, can the entire agent workload be contained?

38North Security's Agentic AI Assurance framework proposes nine augmentation indicators for evaluating those questions in a FedRAMP 20x environment. The indicators are defined by 38North Security—they do not replace FedRAMP KSIs or change the canonical FedRAMP KSI score.

The framework follows four rules:

Every indicator is grounded in published frameworks or federal guidance. Each one traces to requirements or risk-management practices in the NIST AI Risk Management Framework and Generative AI Profile, the OWASP Top 10 for Agentic Applications, or OMB Memorandum M-25-21.

Only observed evidence affects a score. Evidence from running systems can produce a scored PASS or FAIL. Policies, inventories, and other self-declared records remain informational.

Missing evidence does not default to a pass. If neither observed nor attested evidence can satisfy an indicator, the result is a FAIL.

The Blueprint Alliance architecture provides the organizing map. Its four questions help organize the indicators, but Blueprint is not an evidence authority and does not determine results.

How Agentic AI Assurance Fits Into FedRAMP 20x

GROUNDING IN

NIST

AI RMF 1.0 +
Generative AI
Profile

OWASP

Top 10 for Agentic
Applications

OMB

Memorandum
M-25-21

ORGANIZED BY

Blueprint Alliance Architecture

Where are my agents? What can they do?
What are they doing? How do I respond?

Organizing crosswalk, not an evidence authority

38North Agentic AI Assurance

Nine augmentation indicators

INV Agent Inventory	AGY Agent Authority & Scope	TRC Agent Action Trail
EFX Enforced Effect Boundaries	HITL Human Review, Override & Halt	INP Goal & Input Integrity
SCV Agent Supply-Chain Vetting	TST Proportionate Testing & Monitoring	RCV Agent Containment & Recovery

EVIDENCE DETERMINES THE RESULT

Observed evidence

Running-system evidence

→ **PASS / FAIL**

→ contributes to the augmentation-family score

Attested evidence

Customer-maintained governance records

→ **INFORMATIONAL**

→ does not affect the score

No supporting evidence

→ **FAIL**

FedRAMP 20x

Agentic AI Assurance augments the FedRAMP 20x model. It does not replace or modify FedRAMP-defined KSIs or their required validations.

2. Why Agentic AI Requires Additional Assurance

AI agents change the security problem because they do more than process information. They can reason through multi-step tasks, call tools, access data and applications, delegate work, and take actions across interconnected systems.

The Blueprint Alliance describes this shift in *Governing Agentic Execution* as a move from software that follows predefined instructions toward autonomous, multi-step workflows that can decide, act, and connect across systems. The Alliance argues that this autonomy creates new identity, access, delegation, runtime-monitoring, and containment challenges for enterprise security teams.¹

Those questions cannot be answered by configuration alone. Knowing what an agent was provisioned to access does not necessarily establish what it can actually reach or do during execution. The Blueprint Alliance therefore calls for task-scoped access rather than standing privilege, end-to-end traceability of delegation, continuous monitoring of runtime behavior, and the ability to contain agent activity when necessary.²

An agent's own record of its actions is insufficient if it cannot be corroborated by enforcement points and target systems. Human approval must apply to the specific action being authorized, and containment must account for descendant agents and queued work, and not just the agent that initiated the task.

FedRAMP 20x does not yet define agent-specific KSIs for these conditions. 38North Security developed the Agentic AI Assurance framework to address that gap with nine augmentation indicators covering agent inventory, authority and scope, action trails, effect boundaries, human oversight, goal and input integrity, supply-chain vetting, testing and monitoring, and containment and recovery. The indicators remain separate from the canonical FedRAMP KSI score.

3. How the Framework is Built

3.1 Standards and Authorities

Three published authorities provide the foundation for the indicator family:

AUTHORITY	ROLE
NIST AI RMF 1.0 and the Generative AI Profile (NIST AI 600-1)	Risk-management framework. Provides the Govern, Map, Measure, and Manage functions; NIST AI 600-1 adds GAI-specific risk areas: Human-AI Configuration, Information Integrity, Value Chain and Component Integration.
OWASP Top 10 for Agentic Applications (ASI01 to ASI10)	Risk taxonomy. The family maps all ten categories to one or more indicators (Appendix A).
OMB Memorandum M-25-21	Federal AI governance and risk-management requirements for agencies, including use-case inventory, high-impact determinations, pre-deployment testing, ongoing monitoring, and human oversight and intervention.

3.2 The Blueprint Alliance Organizing Model

The Blueprint Alliance is a multi-vendor industry initiative that includes Okta, AWS, CrowdStrike, Databricks, Docker, Google Cloud, Lovable, Proofpoint, Salesforce, ServiceNow, Wiz, and ZScaler. In September 2026, the Alliance published *Governing Agentic Execution: An Architectural Blueprint for the Secure Agentic Enterprise*, an open reference architecture for securing AI agents.

The architecture organizes agent security around four questions:



Agentic AI Assurance uses this architecture as an organizing crosswalk. Each indicator maps to relevant Blueprint questions, architectural components, and principles alongside its NIST, OWASP, and OMB references. This gives the nine indicators a consistent way to describe where they fit within the broader problem of governing agentic systems.

The Blueprint architecture does **not** serve as an evidence authority and does not determine whether an indicator passes or fails. It is an industry reference architecture, not a government standard. The framework’s evidence anchors remain NIST AI 600-1, the OWASP Top 10 for Agentic Applications, and OMB Memorandum M-25-21.

3.3 Terminology: Consequential Use

OMB M-25-21 defines **high-impact AI** as AI whose output serves as a principal basis for decisions or actions with a legal, material, binding, or significant effect on rights or safety.

This framework uses **consequential use** as a 38North Security operational term for agent workflows that warrant stronger controls or oversight. It is not intended to replace or redefine OMB's formal high-impact AI determination.

3.4 Evidence tiers

Each indicator reads evidence through two tiers.

Observed evidence. Evidence read from the running system: cloud audit trails, identity and permission configuration, enforcement-point logs.

Attested evidence. Governance records the customer maintains about each agent, with human provenance: the inventory entry, the consequential-use classification, the oversight policy, the effect-boundary policy, the delegation policy, the tool allowlist, the supply-chain vetting record, test and monitoring records, the recovery procedure, and the record that a principal is a non-agentic use.

Observed evidence is evaluated first. If no observed evidence method is available, an indicator may rely on attested evidence. Attested evidence can support an informational result but never changes the scored percentage.

Section 3.5 explains how evidence provenance determines the result and whether it contributes to the scored mean.

3.5 Scoring Model

Provenance decides scoring, per result, not per indicator. The same indicator is scored for a customer with observed evidence and informational for a customer with only attestation. Uploading a policy document can never move a scored percentage.

WHAT ANSWERED	RESULT	IN A SCORED MEAN?
An observed satisfier	Scored PASS or FAIL	Yes, in the augmentation-family mean
An attested satisfier only	INFORMATIONAL, labeled self-declared, shown with its anchor and the attested record	No
Neither tier reachable	FAIL	Yes, as a failure
A Security Decision Record determination that the customer has no agentic AI in scope	NOT APPLICABLE	No
A “no agentic AI” determination while agent activity is observed	CONFLICT, surfaced for review	Reported, not scored

Three rules complete the model.

- **Multi-leg indicators score only when every leg is proven.** An indicator with several requirements (AGY has three) is scored only when every leg has observed evidence. If any leg is attested only, the whole result is informational, and the provenance of each leg is disclosed. A narrow-observed slice, such as IAM scoping, cannot score a broad claim, such as effect and delegation constraints.
- Every result discloses how it was proven. Each result reports which satisfier answered, whether it was observed or attested, and the evidence source. An assessor sees exactly how each indicator was proven and that an informational result is self-declared.
- The family is separate from the canonical score. Agentic AI Assurance is an augmentation family. Its scored results sit outside the FedRAMP canonical KSI mean. Informational and not-applicable results sit outside any mean and are shown as posture disclosure.

In the first release each indicator discloses one clearly labeled method. A second independent method per indicator, such as cloud-observed scoping that agrees with an attested allowlist, is a later addition.

4. The Nine Agentic AI Assurance Indicators

The indicators are ordered foundation first, then by current observability. Each entry explains what the indicator assesses, its published references, its Blueprint Alliance mapping, and its evidence and scoring model.

The nine indicators, at a glance:

INDICATOR	QUESTION IT ANSWERS
INV Agent inventory	What agents are operating here?
AGY Agent authority and scope	What is each agent allowed to access, delegate, and change?
TRC Agent action trail	Can we reconstruct what the agent actually did?
EFX Enforced effect boundaries	Are consequential effects constrained before they occur?
HITL Human review, override and halt	Where can a human intervene, approve, or stop execution?
INP Goal and input integrity	Can instructions, goals, memory, or inter-agent messages be manipulated?
scv Agent supply-chain vetting	Are models, tools, plugins, and dependencies trusted and revalidated?
TST Proportionate testing and monitoring	Is agent behavior being tested and monitored according to risk?
RCV Agent containment and recovery	Can agent activity be contained and recovered when something goes wrong?

4.1 KSI-AIA-INV: Agent inventory

What it assesses. Whether the organization maintains a current inventory of agentic AI and can reconcile that inventory against observed agent activity.

The customer holds a current inventory of its agentic AI, and each entry carries the agent-specific attributes: the accountable business owner, the authorized purpose, the autonomy level, the effective authority (what the agent can actually reach, not only what it was granted), its dependencies, and its active delegated workloads. A bare list of agents adds little over the automated inventory FedRAMP 20x already requires. The agent-specific attributes are the requirement. Each entry is classified by consequential use.

The inventory is proven current against an observed surface, not by assertion. Every principal observed invoking a model or agent runtime inside the boundary (on AWS, the Bedrock invocation and agent events in CloudTrail) must be either an inventory entry or a recorded non-agentic use. This discovery leg follows the Blueprint Alliance’s treatment of the governed execution path as a detection surface.

ANCHORS	OMB M-25-21 (AI use-case inventory; high-impact determination). NIST AI RMF, Govern function. Canonical FedRAMP relative: KSI-PIY-GIV.
---------	--

BLUEPRINT ALLIANCE MAPPING	■ Question 1, Where are my agents? (Development, Discovery and Identity): the Agent Directory (verified profile, accountable owner, lifecycle status, risk tier, operational scope) and Shadow AI Discovery. ■ Principles: every hosted agent is a distinct security identity; governance adapts at the velocity of AI (a current inventory).
----------------------------	---

EVIDENCE AND SCORING	Governance indicator. The inventory records are attested, so INV never scores PASS on attestation alone. A complete attested inventory reports as informational. The discovery leg is observed and can only lower the result: an unregistered agent principal fails INV, and a "no agentic AI" determination surfaces CONFLICT when invocations are observed.
----------------------	---

4.2 KSI-AIA-AGY: Agent authority and scope

What it assesses. Three legs.

- **Tool and credential scope.** These are least-privilege and explicitly scoped to the agent.
- **Effect constraints.** Permissions bound what the agent can change (the recipient, the record, the state change), not only which tools it may call.
- **Delegation and data constraints.** Policy defines what an agent may delegate, how much authority a child agent receives, and whether another agent's instructions may be treated as commands. Retrieval of protected data is authorized before retrieval. Agent memory is separated per tenant and per task, so data one agent may see cannot flow through persistent memory to an agent that may not.

ANCHORS

OWASP ASI02 (Tool Misuse), ASI03 (Identity and Privilege Abuse), ASI07 (Insecure Inter-Agent Communication). OWASP LLM06 (Excessive Agency).

BLUEPRINT
ALLIANCE MAPPING

- Question 2, What can they do? (Access and Entitlements): fine-grained, time-bound access policies, and delegation and inheritance as the intersection of entitlements (legs 1 and 3).
- Question 3, What are they doing? (Runtime Authorization and Control): Purpose-Aware Data Security (leg 3: data retrieval authorized before retrieval, per-task memory).
- Principles: access is scoped to each task, not standing; delegation is traceable end to end (leg 3).

EVIDENCE AND
SCORING

Probe-strong. The result is scored only when all three legs are observed. Any attested-only leg makes the whole result informational, with each leg's provenance disclosed. This keeps a narrow observed slice (IAM scoping) from scoring the broad claim (effect and delegation constraints).

4.3 KSI-AIA-TRC: Agent action trail

What it assesses. Whether every agent action is logged to an immutable, attributable trail.

The trail is sourced from enforcement points and target systems, never only from the agent’s own account of itself. An immutable record can faithfully preserve a false account.

Each action carries an outcome state: requested, approved, attempted, committed, failed, or unknown. Agent identity passes specific, testable identity-validation checks defined per binding. This replaces the untestable claim “cannot be spoofed”.

ANCHORS

OWASP ASI03 (Identity and Privilege Abuse), ASI10 (Rogue Agents). NIST AI RMF, Measure function. Canonical FedRAMP relative: KSI-MLA-OSM (tamper-resistant centralized logging).

BLUEPRINT
ALLIANCE MAPPING

- Question 3, What are they doing? Monitoring (tracing, continuous monitoring with non-repudiation) and Continuous Intent and Alignment Enforcement (traceable decision evidence).
- Cross-cutting Telemetry, Logging and Observability: “capturing telemetry at the execution boundary rather than relying on self-reported agent logs” (whitepaper p. 25). This is the TRC requirement word for word, and the strongest sign that the two models agree.
- Principles: every hosted agent is a distinct security identity (validated agent identity); delegation is traceable end to end; the agent runtime is isolated and monitored.

EVIDENCE AND
SCORING

Probe-strong. Observed trail evidence from enforcement points and target systems scores. A trail attested by document only is informational.

4.4 KSI-AIA-EFX: Enforced effect boundaries

What it assesses. Whether the business effects an agent can cause are bounded and enforced before they occur.

Checks must occur before a consequential effect, not after it. Every path an agent can use to create that effect —API, browser, shell, code execution, or another route—must pass through an enforcement point. Blocking one route is not enough if the agent can achieve the same outcome through an alternate path, such as a CI/CD restriction bypassed through a leftover sudoers grant.

Authorization must also account for accumulated effects. A series of individually permitted actions can combine into a prohibited outcome, so enforcement must consider the state of the task as it progresses across tools, time, and other agents.

ANCHORS

OWASP ASI05 (Unexpected Code Execution), ASI02 (Tool Misuse). OWASP LLM06 (Excessive Agency). OMB M-25-21 (minimum risk-management practices). NIST AI RMF, Manage function.

- BLUEPRINT ALLIANCE MAPPING
- Question 2, What can they do?: guardrails and effect ceilings.
 - Question 3, What are they doing?: Gateways as policy enforcement points on every route (AI/LLM, MCP, security, agent, API), and Continuous Intent and Alignment Enforcement (plan versus action evaluation, the stateful task constraint).
 - Principle: access is scoped to each task, not standing.

EVIDENCE AND SCORING

Probe or attest. Observed enforcement-point evidence scores. Policy attestation alone is informational.

4.5 KSI-AIA-HITL: Human review, override and halt

What it assesses. Whether agents in consequential uses have defined human review, override, and halt.

Where approval is required, the approval binds to the exact recipient, amount, payload, and relevant state. It expires and cannot be reused.

A halt stops the whole workload: descendant agents and queued work, not only the planner.

Human oversight is also tested. Reviewer effectiveness is evaluated using seeded errors under realistic workloads.

ANCHORS	OMB M-25-21 (human oversight and intervention). NIST Generative AI Profile, "Human-AI Configuration". OWASP ASI09 (Human-Agent Trust Exploitation).
BLUEPRINT ALLIANCE MAPPING	■ Question 2, What can they do?: Governance (contingent review). ■ Question 3, What are they doing?: Monitoring (workflow-in-the-loop). ■ Question 4, How do I respond? (Active Containment): a halt stops the whole workload. ■ Principle: containment is instant and reversible.
EVIDENCE AND SCORING	Probe or attest. Observed approval-binding and halt evidence scores. An oversight policy alone is informational.

4.6 KSI-AIA-INP: Goal and input integrity

What it assesses. Whether the system protects agent goals and inputs from manipulation, including goal hijacking, prompt injection, memory poisoning, and unauthorized inter-agent instructions.

Before execution, the system validates that the agent’s goal is authorized and correctly understood. Output from another agent is not automatically trusted as an authorized command.

ANCHORS	OWASP ASI01 (Agent Goal Hijack), ASI06 (Memory and Context Poisoning), ASI07 (Insecure Inter-Agent Communication). NIST Generative AI Profile, "Information Integrity".
---------	---

BLUEPRINT ALLIANCE MAPPING	■ Question 3, What are they doing?: Monitoring (prompt-injection inspection) and Continuous Intent and Alignment Enforcement (purpose binding, plan versus action evaluation). ■ Principle: the agent runtime is isolated and monitored, not just provisioned.
----------------------------	--

EVIDENCE AND SCORING	Attest, with an observed probe rare. Expect informational results until runtime inspection evidence is observable. The attested defenses and the goal-validation policy are disclosed with each result.
----------------------	---

4.7 KSI-AIA-SCV: Agent supply-chain vetting

What it assesses. Whether third-party tools, plugins, models, and agents are vetted and integrity-checked.

Vetting is repeated after a material change to the agent, including changes to its model, prompts, tools, permissions, orchestration, memory configuration, or internally generated tools.

The customer defines what constitutes a material agent change.

ANCHORS

OWASP ASI04 (Agentic Supply Chain Vulnerabilities). NIST Generative AI Profile, "Value Chain and Component Integration". Canonical FedRAMP relative: KSI-CMT-VTD (validation throughout deployment).

- BLUEPRINT
ALLIANCE MAPPING
- Question 1, Where are my agents?: AI Agent Security Posture Management (code and image provenance, tool and skill supply-chain scanning).
 - Question 4, How do I respond?: Access Recovery (re-attestation after a material change).
 - Principle: governance adapts at the velocity of AI (revalidation after material change).

EVIDENCE AND
SCORING

Probe or attest. Observed vetting and integrity-check evidence scores. A vetting record alone is informational.

4.8 KSI-AIA-TST: Proportionate testing and monitoring

What it assesses. Whether all agents get testing and monitoring proportionate to their use, stronger for consequential uses.

The test set covers ordinary mistakes, bypass paths, concurrent actions, failed policy services, stale approvals, and unsuccessful revocation, not only prompt injection.

ANCHORS OMB M-25-21 (pre-deployment testing; ongoing monitoring). NIST AI RMF, Measure and Manage functions.

- BLUEPRINT ALLIANCE MAPPING**
- Question 1, Where are my agents?: AI Agent Security Posture Management (adversarial validation).
 - Principles: the agent runtime is isolated and monitored; governance adapts at the velocity of AI.

EVIDENCE AND SCORING Governance indicator. In the initial implementation, TST relies on customer testing and monitoring records rather than an independently observed probe. Results are therefore informational and do not affect the scored mean.

4.9 KSI-AIA-RCV: Agent containment and recovery

What it assesses. Whether the organization can contain an agentic workload and recover from failure by revoking outstanding authority, terminating descendant agents and queued work, and addressing cascading effects.

Recovery capability is demonstrated through a drill, not merely documented in a procedure.

ANCHORS

OWASP ASI08 (Cascading Failures), ASI10 (Rogue Agents). NIST AI RMF, Manage function.

- BLUEPRINT ALLIANCE MAPPING
- Question 4, How do I respond? (Active Containment): Enforcement actions (token and session revocation, throttling, platform termination, kill-switch event protocol) and Access Recovery (re-attestation, staged re-enrollment, root-cause sign-off, audit-trail closure).
 - Cross-cutting Execution Context and Risk Signals: a risk signal drives containment.
 - Principle: containment is instant and reversible.

EVIDENCE AND SCORING

Probe or attest. A drill observed in identity and runtime audit events scores. A recovery procedure without a drill is informational.

5. Implementation Status

The Agentic AI Assurance indicator family defines the assurance model described in this paper.

NorthWatch implementation is in progress. At publication, the complete indicator set and scan results described here should not be interpreted as generally available production capabilities.

6. Conclusion

Agentic systems create assurance problems that cannot be addressed through documentation alone. Their authority, actions, delegation, effects, and recovery capabilities must increasingly be evaluated through evidence produced by the systems in which they operate.

The Agentic AI Assurance framework is an initial model for doing that within a FedRAMP 20x environment: grounding each indicator in published authority, distinguishing observed evidence from self-declaration, and keeping the resulting augmentation separate from FedRAMP's canonical KSI score.

As agentic architectures and federal guidance evolve, the framework can evolve with them. The underlying principle should remain stable:

Assurance should reflect what the system can demonstrate, not merely what it claims.

Appendix A: OWASP Coverage

OWASP Top 10 for Agentic Applications

OWASP CATEGORY	TREATED BY
ASI01 Agent Goal Hijack	INP
ASI02 Tool Misuse	AGY EFX
ASI03 Identity and Privilege Abuse	AGY TRC
ASI04 Agentic Supply Chain Vulnerabilities	SCV
ASI05 Unexpected Code Execution	EFX
ASI06 Memory and Context Poisoning	INP
ASI07 Insecure Inter-Agent Communication	AGY (delegation and data constraints), INP (instruction handling)
ASI08 Cascading Failures	RCV , EFX (accumulated effects)
ASI09 Human-Agent Trust Exploitation	HITL
ASI10 Rogue Agents	TRC RCV

Appendix B: Blueprint Alliance Crosswalk

By question

BLUEPRINT QUESTION (PILLAR)	INDICATORS
Where are my agents? (Development, Discovery and Identity)	INVSCVTST
What can they do? (Access and Entitlements)	AGYEFXHITL
What are they doing? (Runtime Authorization and Control)	EFXINPTRCHITLAGY
How do I respond? (Active Containment)	RCVSCVHITL
Cross-cutting: Execution Context and Risk Signals; Telemetry, Logging and Observability	TRCRCV

By principle

BLUEPRINT PRINCIPLE	INDICATORS
Every hosted agent is a distinct security identity, not an afterthought	INVTRC
Access is scoped to each task, not standing	AGYEFX
Delegation is traceable end to end	AGYTRC
Agent runtime is isolated and monitored, not just provisioned	TSTTRCINP
Containment is instant and reversible	RCVHITL
Governance adapts at the velocity of AI	INVSCVTST

References

- NIST AI 100-1, *Artificial Intelligence Risk Management Framework* (AI RMF 1.0, 26 January 2023).
- NIST AI 600-1, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile* (26 July 2024).
- *OWASP Top 10 for Agentic Applications* (ASI01 to ASI10), and the *OWASP Top 10 for LLM Applications* (LLM06, Excessive Agency, 09 December 2025).
- OMB Memorandum M-25-21, *Accelerating Federal Use of AI through Innovation, Governance, and Public Trust* (03 April 2025).
- Blueprint Alliance, *Governing Agentic Execution: An architectural blueprint for the secure agentic enterprise* (22 September 2026, blueprintalliance.ai).
- FedRAMP 20x Key Security Indicators (canonical relatives KSI-PIY-GIV, KSI-MLA-OSM, KSI-CMT-VTD).

-
1. Blueprint Alliance, *Governing Agentic Execution: An Architectural Blueprint for the Secure Agentic Enterprise*, 22 September 2026.
 2. Blueprint Alliance, *Governing Agentic Execution*, 22 September 2026.

About the Author



AUTHOR

Austen Summers

Principal Architect, [NorthWatch](#)

I'm a cybersecurity engineer and software architect focused on federal cloud security and compliance. My experience ranges from assessing cloud environments for the FAA to building the software that automates much of that work.

At 38North Security, I led the team that designed and built [NorthWatch](#), our FedRAMP 20x compliance suite. It helps cloud service providers earn and maintain the authorization they need to sell to the federal government by bringing continuous evidence collection, automated control validation, and vulnerability evaluation into a single platform. Providers can run it hosted by us or inside their own cloud. I also built Arcanum, our internal AI platform for working with controlled client data. Beyond the software, I developed our vulnerability risk scoring model for FedRAMP 20x and an overlay that extends its Key Security Indicators to agentic AI.

Before 38North, I spent three years at the Georgia Tech Research Institute, advancing from research assistant to secure computing engineer. I reviewed systems against NIST standards, helped design and operate GTRI's enterprise WAN, built DevSecOps pipelines, hardened research infrastructure, and automated deployments so nobody had to do things the hard way twice.

I've written on FedRAMP 20x, released open-source deployment tools, and mentored students through Kennesaw State University's Cybersecurity Capstone Mentorship program. I do my best work where the problems are complex, the stakes are real, and there's always something new to figure out.