HUMAN

Agentic AI in Government:

Balancing Autonomy and Accountability

As autonomous AI systems take on more complex roles, government leaders must weigh the promise of faster decisions against the need for control, oversight and accountability.

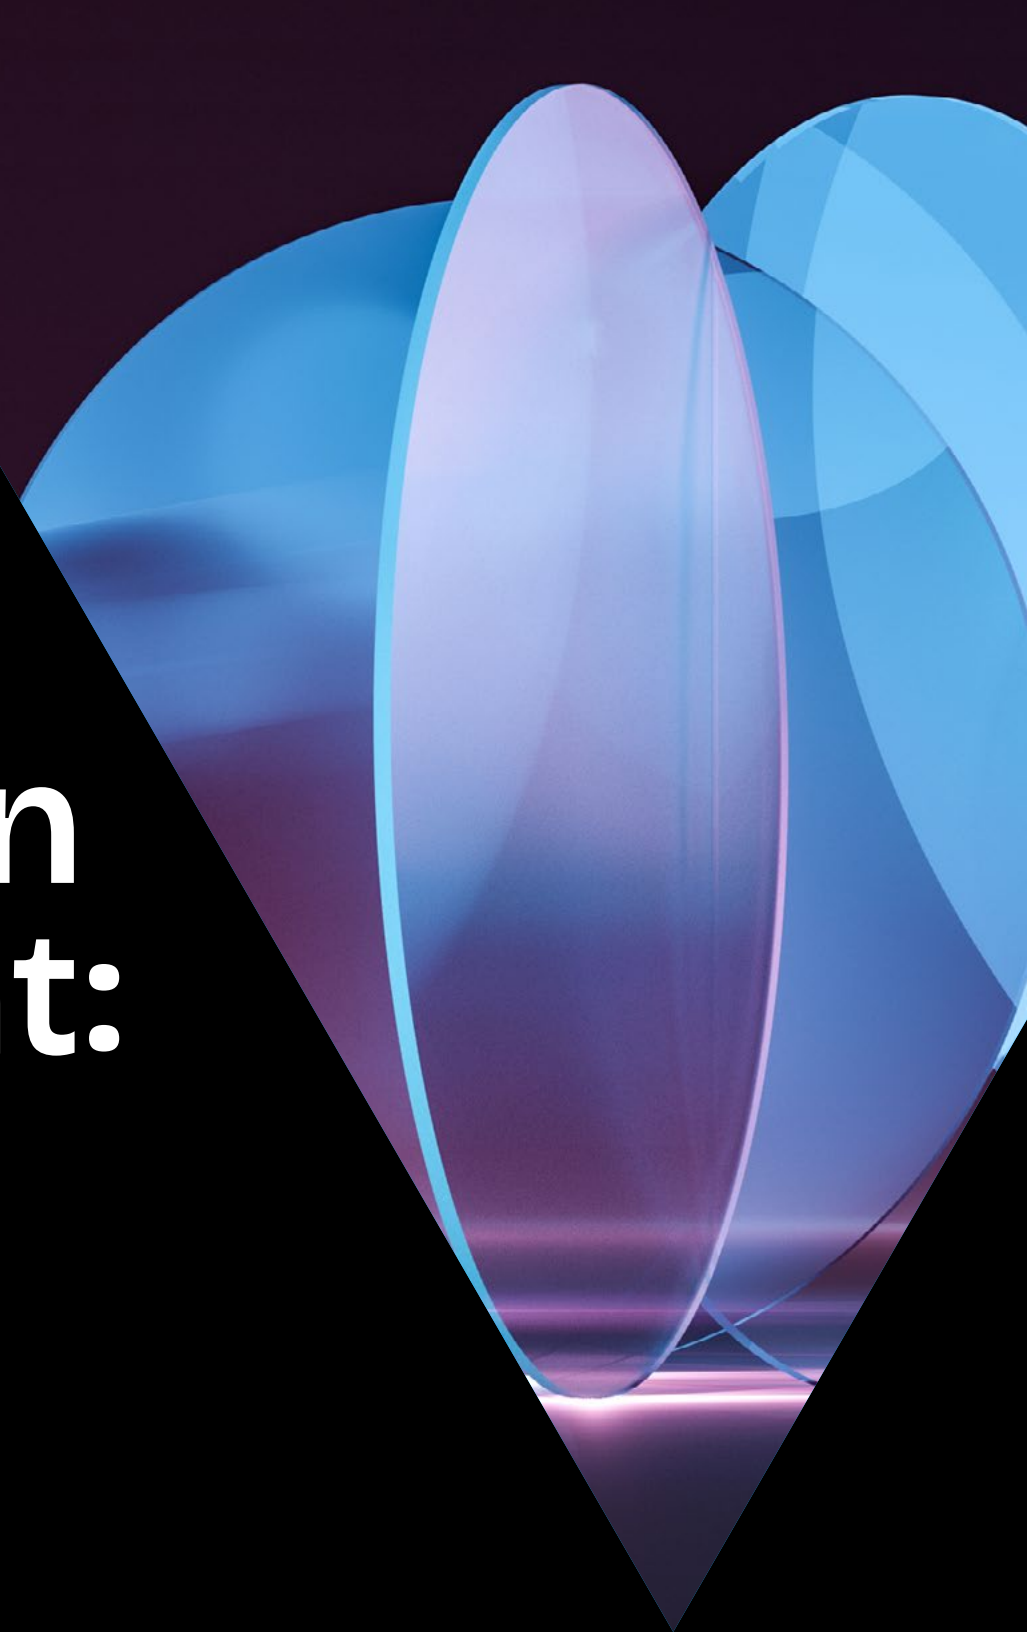

Introduction

Artificial intelligence is already transforming how government agencies process data, respond to threats and deliver services. But agentic AI — systems that can make autonomous decisions, perform complex tasks and interact across platforms independently — signals a deeper, more disruptive shift.

Unlike traditional systems that simply analyze and inform, agentic AI can act. It doesn't just offer recommendations; it carries them out, making decisions, initiating actions and learning from outcomes with minimal human input.

This leap from reactive to proactive AI brings enormous potential and equally significant risks. Agentic systems can reason, remember and access external tools, enabling them to function more like autonomous teammates than passive assistants. That ability could accelerate decision-making across government, from battlefield scenarios to back-office operations — but it also raises pressing questions about oversight, trust and accountability.

These questions were front and center during a recent GovExec webcast, **Decoding Agentic AI**, as Captain Christopher Clark, artificial intelligence lead for the U.S. Marine Corps and Frank Walsh, HUMAN's field chief technology officer, explored how this new class of AI is reshaping the cyber landscape. While defenders are gaining tools to rapidly triage threats and probe vulnerabilities, attackers are also becoming more sophisticated. In this emerging arms race, both sides are leveraging agentic systems, prompting agencies to think carefully about how to deploy these technologies safely and strategically.

Table of Contents

HUMAN

01.

Agentic AI is a shift in both cyber attacks and defense

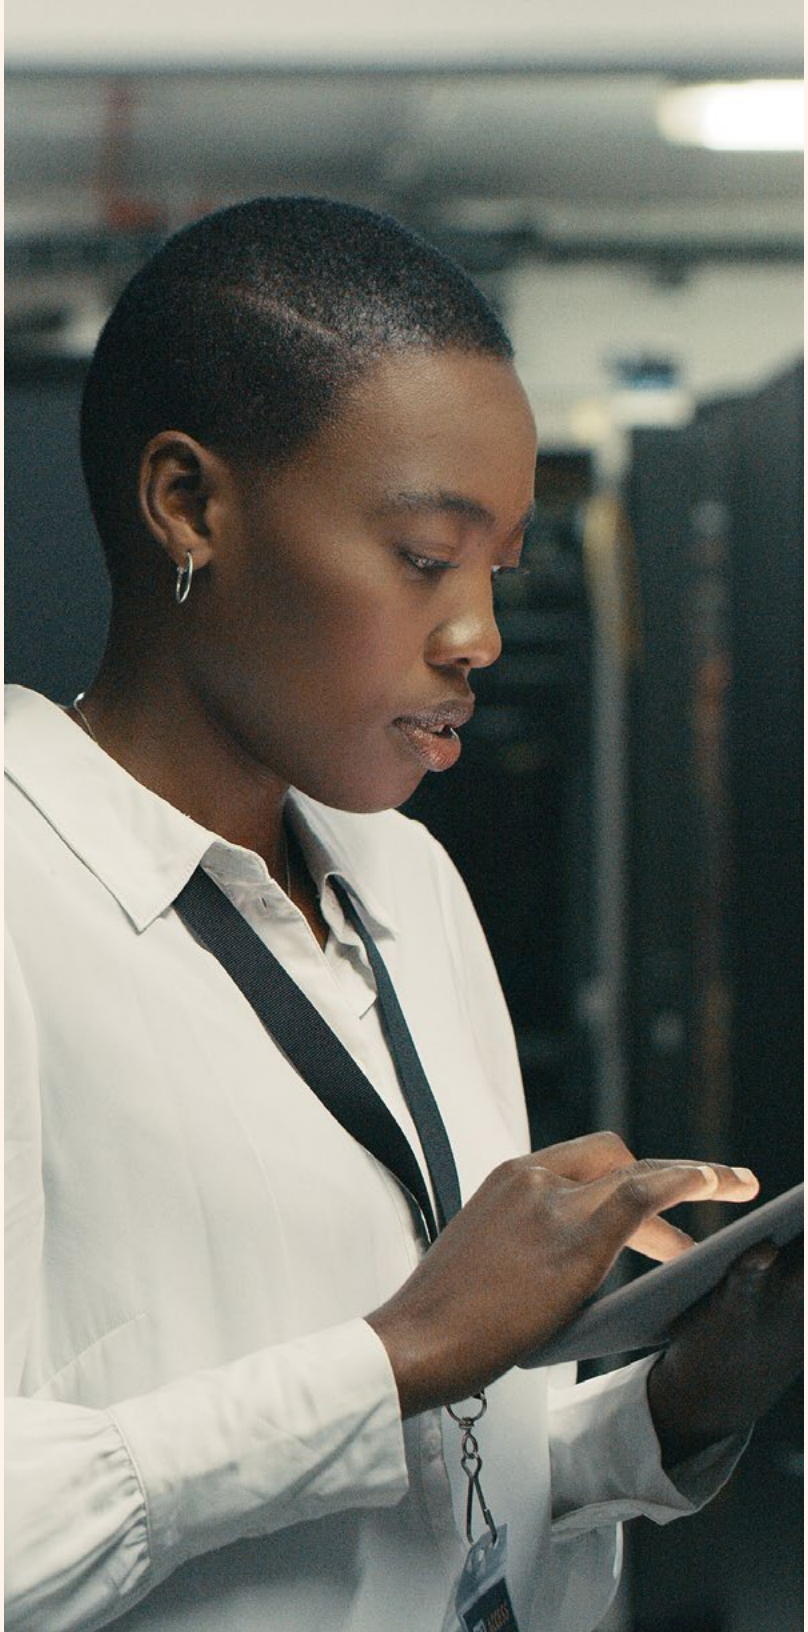

AGENTIC AI REPRESENTS MORE than just an evolution of traditional automation — it's a fundamental shift in how both attackers and defenders operate in the cyber domain.

On the offensive side, Walsh outlined the advantages threat actors are gaining: "I like to use the alliteration: speed, scale, stealth, and script, and in all four of those capacities, we're seeing agentic artificial intelligence offering attackers significant improvements in their capabilities."

According to Walsh, these systems allow adversaries to rapidly iterate on payloads, probe applications in real time, and even generate malicious tooling with AI's help. "They're able to scale these attacks across every input on every interface, continuously using agentic AI to essentially do their bidding."

Yet, agentic AI offers promise for defenders, too, if deployed wisely. "I see tremendous advantages across the board, whether it be characterizing, classifying, ranking the risk of specific threats and offering, 24/7, 365, a way to take action on threats," said Walsh. " We're seeing ways in which we can triage and prioritize the more important attacks. We can penetration test our own software and find gaps, or basically perform static and dynamic code analysis with something far more intelligent."

These agentic systems don't just act; they interact with other services, APIs, and content repositories. In the future, governments may also need to manage how agents access external data, using controls that meter usage, validate authenticity, and prevent unauthorized scraping or misuse. This will be especially important in areas like intelligence, research, and decision-making workflows.

> **I see tremendous advantages across the board, whether it be characterizing, classifying, ranking the risk of specific threats and offering, 24/7, 365, a way to take action on threats."**

FRANK WALSH
Field Chief Technology Officer, HUMAN Security

This evolution has sweeping implications across domains. "I don't think there will be any domain spared as we move from where we are right now into agentic AI," said Capt. Clark. "We're very much interested in how we can inject AI systems into the decision-making process so they can act faster than the enemy."

In short, agentic AI is already altering the tactical landscape. But governments and organizations must balance agility and accountability as they adopt these rapidly evolving technologies.

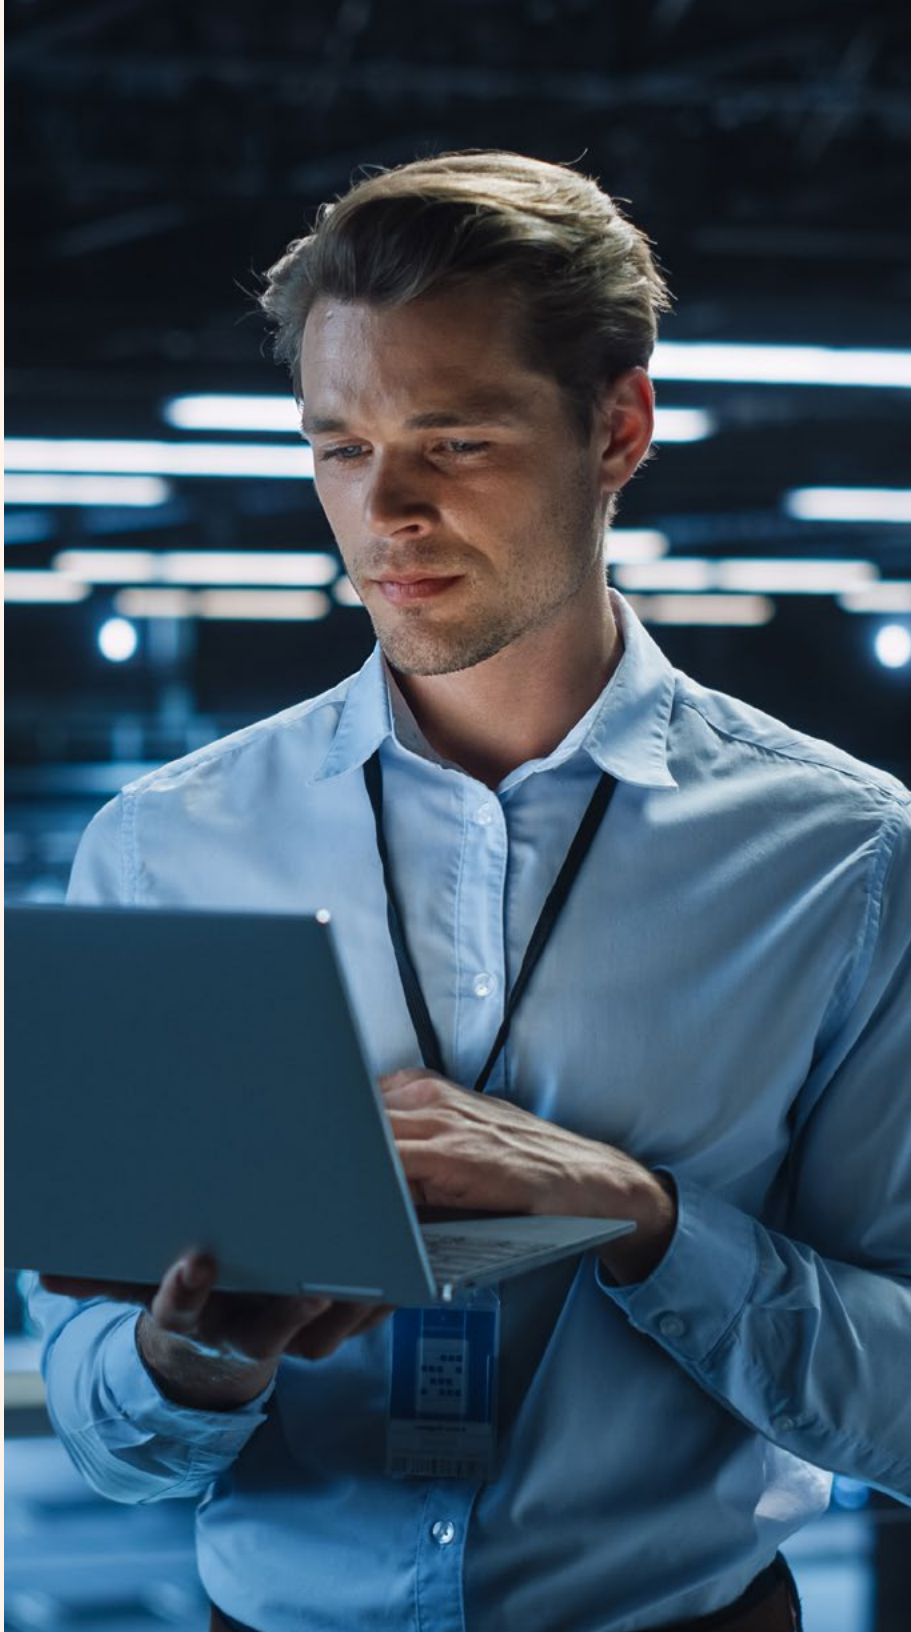

HUMAN

02.

Human oversight and accountability must be built in

AS AGENCIES BEGIN INTEGRATING AGENTIC AI into operations, quality control and clear accountability are paramount. Walsh emphasized that today's AI models, while powerful, still demand close supervision. Whether generating code or drafting decisions, these systems can produce flawed or incomplete outputs if left unchecked.

To that end, the Marine Corps is building a comprehensive governance framework focused on oversight, transparency and explainability. Every deployment of agentic AI must include a human "in the loop" — able to monitor, audit and intervene if the system behaves in unexpected or inappropriate ways.

In fact, the risks of agentic AI are already being cataloged by the security community. The OWASP Foundation recently released a list of 15 distinct agentic AI threat categories — including memory poisoning, tool misuse, and identity spoofing — each of which underscores the need for strong oversight and well-defined controls. Read our breakdown of the OWASP threats on the HUMAN blog.

Walsh echoed the need for continuous oversight, pointing out that agentic AI is not static software — it's a dynamic system that evolves over time. "As we adopt these solutions, we want to understand how our oversight and contribution, or our training to refine these systems, will lead to a more experienced system that persists the lessons we teach it," he said.

For these systems to function like trusted team members, agencies must actively train, monitor and guide them to ensure they internalize the right institutional knowledge. Agencies must therefore develop, test and refine their own governance models that can evolve alongside the technology and remain responsive to new risks as they emerge.

> **"As we adopt these solutions, we want to understand how our oversight and contribution, or our training to refine these systems, will lead to a more experienced system that persists the lessons we teach it."**
>
> **FRANK WALSH**
> Field Chief Technology Officer, HUMAN Security

HUMAN

03.

Start small, work big

AS GOVERNMENT AGENCIES EXPLORE AGENTIC AI, Walsh emphasized the importance of starting small with low-risk, manageable use cases and working up to more complex deployments. This measured approach allows organizations to experiment safely, build institutional knowledge and ensure the right controls are in place before scaling.

"Start small, work big," said Capt. Clark. "I think it's important to have a vision. We may be looking out five, 10 years from now, but being prepared for that is important."

In the Marine Corps, that preparation begins with administrative and business tasks that are clearly low-risk. "Things like language translation or writing documents, reviewing documents, those kinds of things are low risk," Capt. Clark said. These areas provide a foundation for experimentation while maintaining a strong degree of human control and accountability. "These systems cannot be held responsible. We can only hold humans responsible."

Ultimately, beginning with low-risk use cases gives agencies the space to learn, adapt, and put the right guardrails in place before deploying agentic AI in more sensitive or mission-critical contexts.

The aim isn't rapid replacement of human judgment, but thoughtful integration of AI as a force multiplier. By moving deliberately and maintaining clear lines of accountability, government leaders can lay the groundwork for agentic systems that enhance rather than compromise trust, performance, and public service.

> **"Start small, work big. I think it's important to have a vision. We may be looking out five, ten years from now, but being prepared for that is important."**
>
> **CAPTAIN CHRISTOPHER CLARK**
> Artificial Intelligence Lead, U.S. Marine Corps

HUMAN

Watch the full webinar here and explore how HUMAN is helping agencies defend critical infrastructure against evolving threats.