Recorded Future®

EXECUTIVE INSIGHTS REPORT | MAY 2026

Hacking Embodied AI

From Insikt Group®

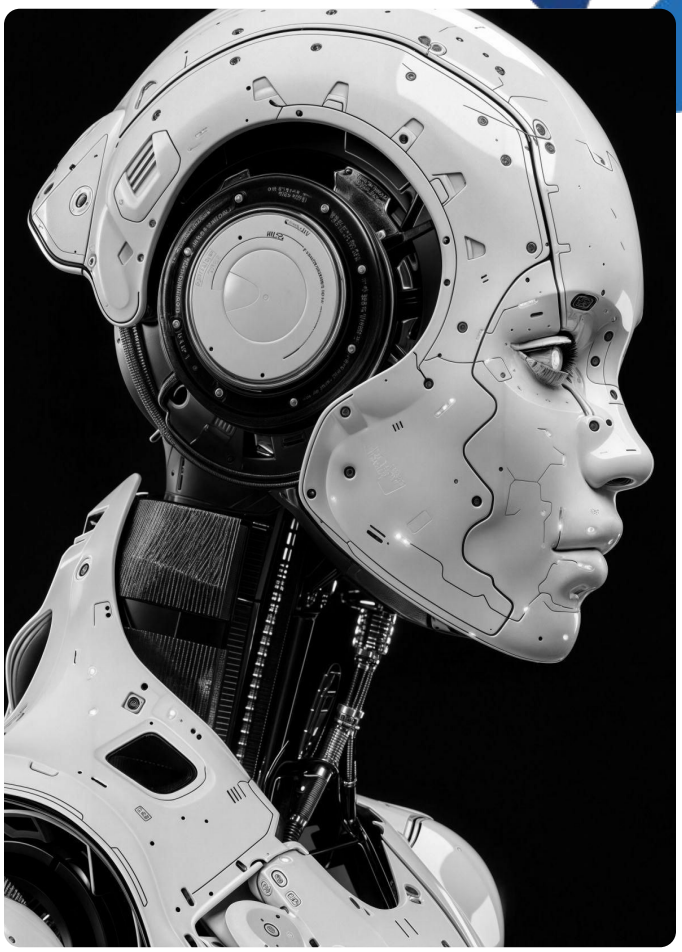

Summary

Embodied AI has arrived. Humanoid and quadruped robots are moving off factory floors and into everyday operations, military deployments, and critical infrastructure. Technological advances in large language models (LLMs) and robotics are enabling robots to perform complex tasks autonomously.

Security has not kept pace. Researchers have demonstrated that commercially available robots can be hijacked over Bluetooth, covertly exfiltrate audio, video, and spatial data to servers in China, and even infect neighboring robots wirelessly, forming physical botnets. If unaddressed, these security weaknesses are set to scale massively once humanoid robots are fully integrated into critical workflows.

The risks need to be taken extremely seriously. A robot should be treated less like a machine on the balance sheet and more like a cyber-physical endpoint with cameras, microphones, radios, cloud dependencies, and motors. That means tougher procurement, tighter network controls, continuous vulnerability monitoring, and a credible plan for operational continuity if a fleet has to be pulled offline.

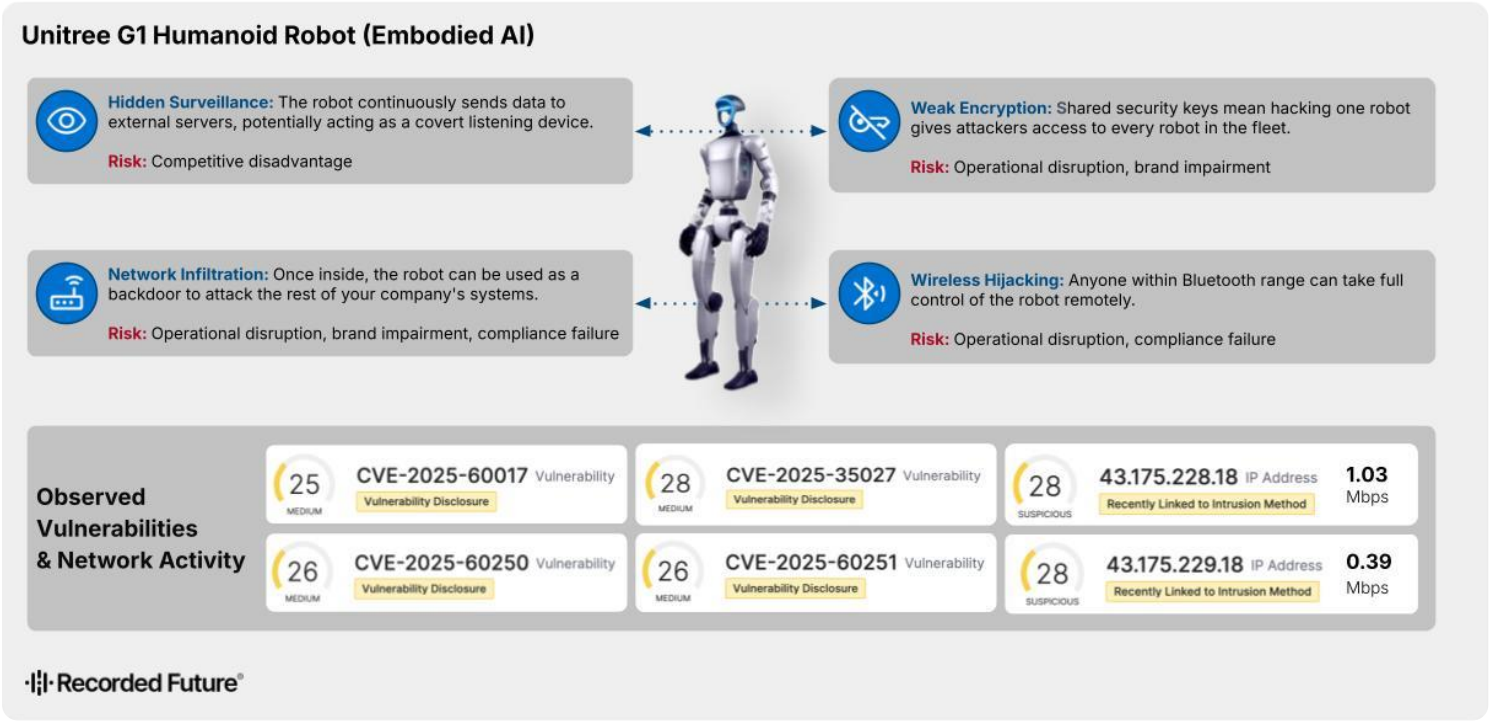

Figure 1: Summary of Unitree G1 vulnerabilities, associated business risks, mapped CVEs, and observed network activity (IPs and data exfiltration rates) (Source: Recorded Future)

Analysis

Market Drivers of Embodied AI Adoption

Embodied AI, intelligent systems in physical forms such as humanoid and quadruped robots, is moving from spectacle to staffing plans.

The shift is being driven as much by demographics as by technological progress. There are growing reports that the working-age population worldwide has begun to decline. China, an economic success story, has seen its population also decline again in 2025 as births hit a record low. These trends do not make large-scale automation inevitable, but they seriously strengthen the economic case for it in both corporate and government decision-making.

The International Federation of Robotics identifies labor shortages, real-world testing of humanoid robots, and increasing attention to safety and cybersecurity as defining trends for 2026. Some early deployments of embodied AI reinforce this trajectory. BMW reports that the Figure 02 humanoid robot has assisted in the production of more than 30,000 X3 vehicles, while GXO and Agility Robotics describe their partnership (established in 2024) as "the first formal commercial deployment of humanoid robots." In high-risk environments, Sellafield is deploying quadruped robots to reduce human exposure in nuclear decommissioning.

Capital markets are also responding. Unitree filed for a reported $610 million initial public offering (IPO) in Shanghai in March 2026. Taken together, these signals suggest that robots are leaving pilot programs and becoming operational. That transition makes the security question immediate rather than theoretical.

Expanding Attack Surface in Embodied AI Systems

Unlike traditional IT assets, embodied AI systems combine multiple high-risk components in a single platform: cameras, microphones, sensors, wireless radios, cloud connectivity, and physical actuation. This convergence creates a broad and under-secured attack surface.

A compromised robot can exfiltrate sensitive environmental and operational data, provide persistent remote access to internal networks, and interact physically with its environment, potentially causing unintended physical effects. This elevates robots from conventional endpoints to cyber-physical systems with both digital and real-world consequences.

The risk is compounded by architectural choices. Many platforms rely on cloud-dependent telemetry, wireless provisioning interfaces, and centralized control mechanisms. These design decisions create multiple entry points for attackers and increase the likelihood of compromise across entire fleets of embodied AI systems.

Demonstrated Vulnerabilities and Exploits

The risks are no longer theoretical. Documented vulnerabilities show that commercially available robots can be compromised with relative ease. Unlike traditional cyber threats, which mostly affect the digital world, exploiting robots enables attackers to manipulate the physical world, maximizing the potential for harm.

In 2025, researchers discovered an undocumented backdoor in Unitree's Go1 quadruped robot that enabled remote access via the CloudSail service. Axios reported that an exposed web application programming interface (API) could allow attackers to locate devices globally and, if a robot was online, view live camera feeds without authentication. Where default credentials remained unchanged, full device control was possible. Whether described as a backdoor or a design failure, the implication is the same: robots may be reachable in ways operators do not anticipate, just like any other Internet of Things (IoT) device.

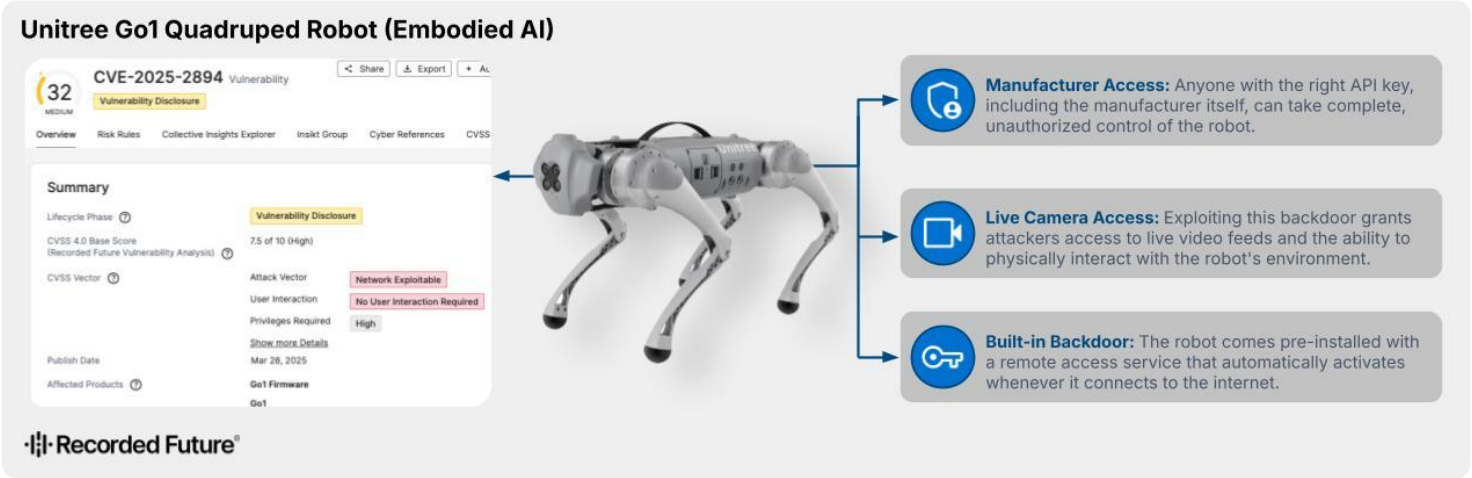

Figure 2: *Summary of vulnerabilities affecting the Unitree Go1 robot, with Intelligence Card insights from the Recorded Future Intelligence Operations Platform (Source: Recorded Future)*

Further research disclosed a critical vulnerability in the Bluetooth Low Energy and Wi-Fi provisioning interface used by multiple Unitree models, including the Go2, B2, G1, R1, and H1 robots. According to both the UniPwn research and IEEE Spectrum, the flaw combined hard-coded cryptographic keys, trivial authentication bypass, and command injection in the Wi-Fi setup process. An attacker within radio range could obtain root-level access without physical contact, giving them control over the robot.

Because the exploit propagates wirelessly, a single compromised device can enable lateral movement across nearby robots. This creates a fleet-level compromise scenario in which multiple units can be controlled simultaneously. The result resembles a physical botnet capable of both digital and physical actions.

Surveillance risks are equally significant. Researchers wrote that the Unitree G1 robot continuously exfiltrated multimodal sensor and service-state telemetry every 300 seconds without the operator's knowledge. This included streaming data to external servers, potentially including audio, video, and spatial mapping. A robot operating inside a plant or laboratory may therefore be mapping the environment in real time.

The attack surface extends beyond firmware and networking layers. Researchers showed they could take control of a Unitree humanoid in about a minute, bypass its normal controller, and trigger physical actions. Demonstrations at GEEKCon in Shanghai indicated that both voice commands and short-range wireless exploits could hijack robots and propagate attacks to nearby units, including those not actively in use.

At the software layer, embodied AI systems introduce additional risks due to their reliance on large vision-language models. Researchers demonstrated that physical-world text can influence system behavior, as injected visual prompts were shown to steer autonomous driving, drone landing, and tracking tasks without compromising the underlying software. This would enable threat actors to take control of a self-driving car or turn a drone into their own surveillance feed by embedding a visual prompt in the environment, such as hiding a message on a stop sign.

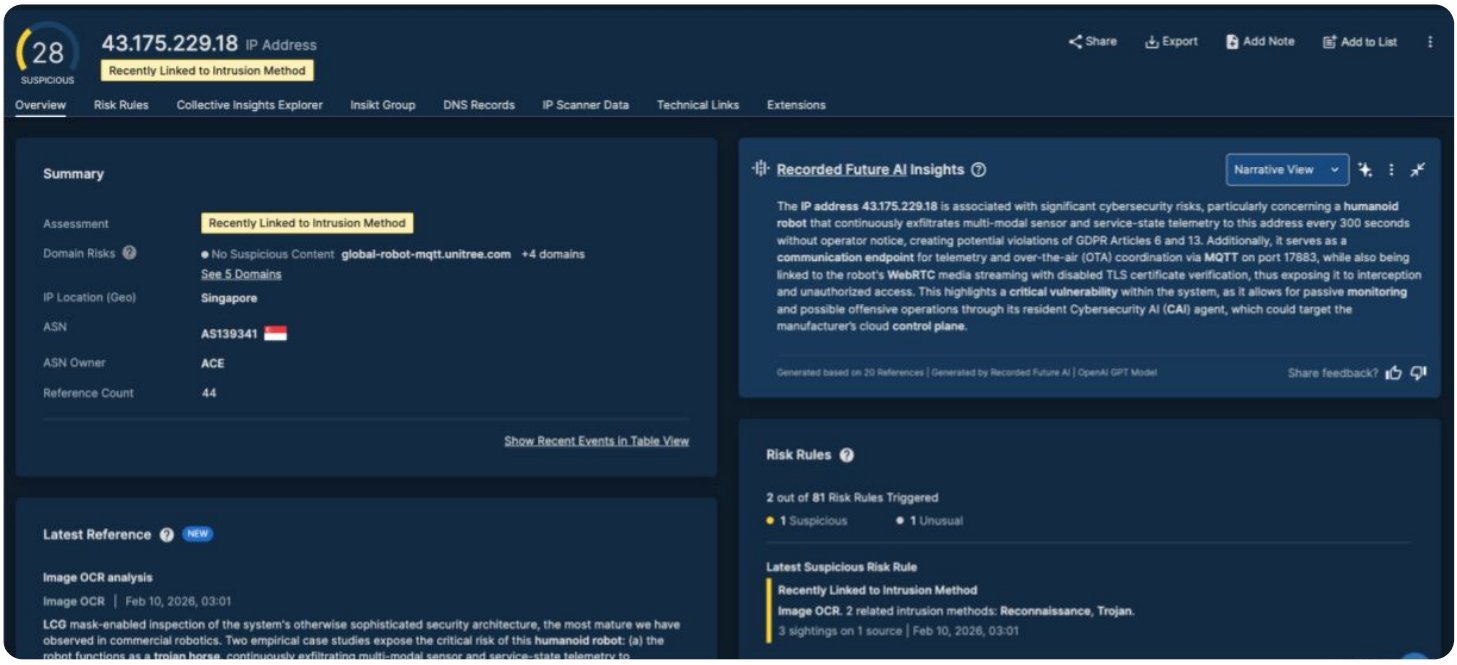

Figure 3: *Researchers found Unitree's G1 quietly transmitting audio, video, and sensor data to the IP address (43[.]175[.]229[.]18) without user awareness (Source: Recorded Future)*

Systemic and Operational Risk Implications

The implications extend beyond individual devices to organizational and systemic risk. Embodied AI systems are already being deployed in environments where compromise has consequences beyond data loss. Manipulation or malfunction of robots during critical operations would have outsized economic or public safety consequences. Militaries are also experimenting with robotic systems (see **Figure 4**).

Figure 4: *Chinese robotic systems demonstrated during military training exercises (left) (Source: ABC YouTube); Concept rendering of the Atlas 2.0 robot operating in a next-generation factory environment (right) (Source: Boston Dynamics YouTube)*

In 2024, the Golden Dragon exercise between Cambodia and China featured robot dogs among the systems on display. Meanwhile, in the US, politicians have begun pushing for Unitree to be designated as a federal supply-chain risk, reflecting national security concerns about commercial robotics platforms.

This is a very similar move to Poland's ban on sensor-rich vehicles accessing military sites to limit surveillance risk. Ukraine has successfully deployed ground-based robots and drones in combat operations, marking a significant shift in modern warfare. In a landmark operation in April 2026, Ukrainian forces captured a Russian position using only unmanned systems — the first recorded instance of a robot-only assault in the conflict.

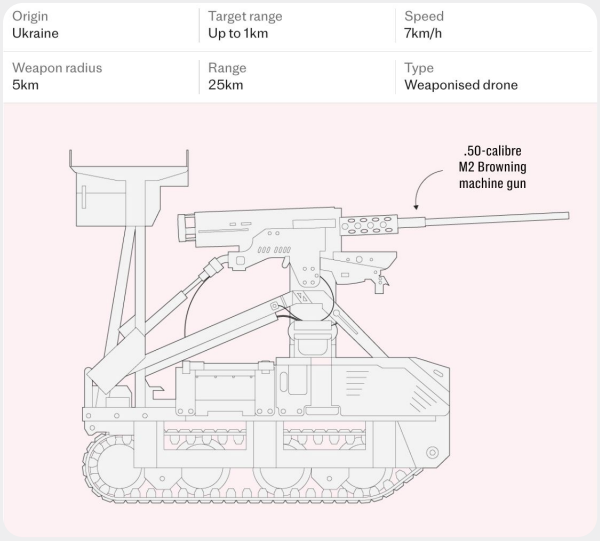

Figure 5: *Droid TW 12.7 machine gun drone, deployed by Ukrainian forces to capture Russian positions without ground troops (Source: The Telegraph)*

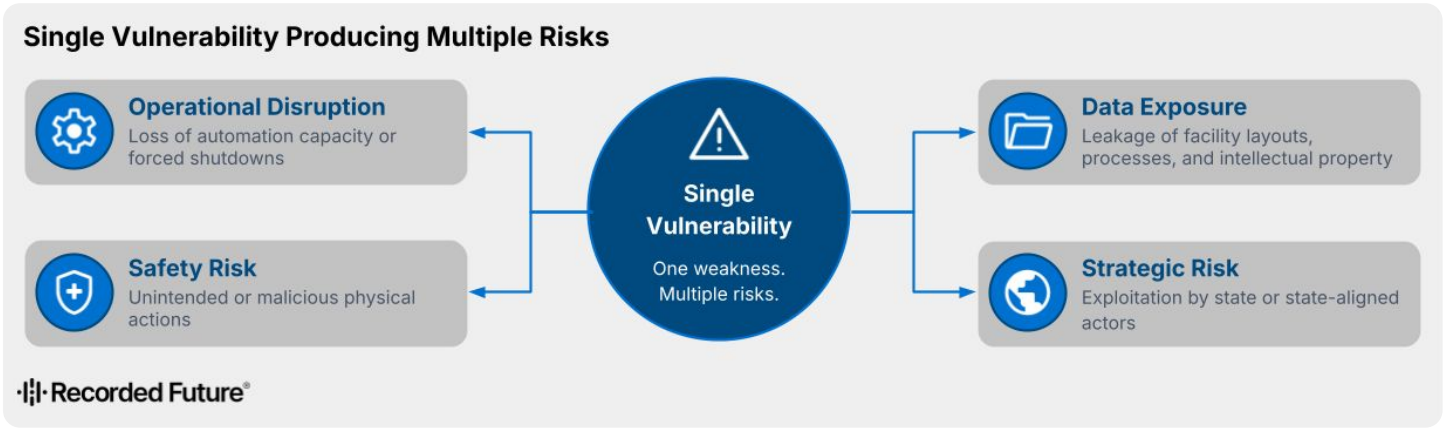

Figure 6: *A single vulnerability can simultaneously produce operational, data, safety, and strategic risks (Source: Recorded Future)*

As adoption scales, these risks become interconnected. A vulnerability affecting one platform or vendor could propagate across fleets, sites, or sectors, creating systemic exposure.

At the same time, the pace of commercial development is outstripping regulatory oversight. Bank of America estimates that as many as three billion humanoid robots could be in operation by 2060. This convergence of demographic pressure, advancing AI capabilities, and falling production costs suggests that large-scale human-machine coexistence is increasingly probable.

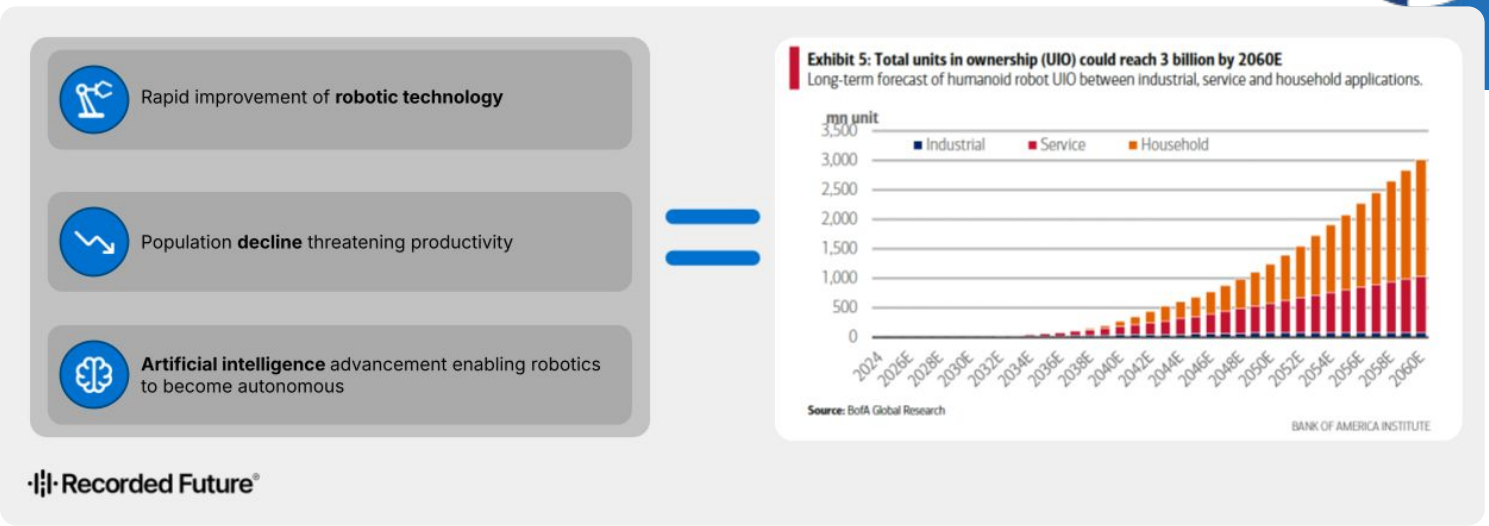

Figure 7: Summary of the factors fueling growth in robotics production, illustrated by Bank of America data (Source: Recorded Future)

Securing embodied AI systems is therefore not a peripheral technical issue. It is a strategic requirement that must be addressed before widespread deployment locks in insecure architectures at scale.

Organizational Risk	Risk Level	How?
Operational disruption	High	Compromised robots halt production, trigger safety shutdowns, or require the full fleet to be removed during incident response.
Competitive disadvantage	High	Exfiltration of proprietary processes, facility layouts, or manufacturing data enables competitors or state-backed firms to replicate them.
Brand impairment	Medium	Publicized incidents involving hijacked or malfunctioning robots erode trust, particularly where safety was a stated benefit.
Financial fraud	High	Compromised robots used as network footholds enable lateral movement to financial systems or support business email compromise and transaction manipulation.
Legal or compliance failure	Medium	Unauthorized data transfer, surveillance, or safety incidents trigger regulatory investigations, fines, and contractual breaches.

Table 1: Business risks associated with the adoption of insecure embodied AI systems (Source: Recorded Future)

Outlook

The mass-production surge: Established car manufacturers and tech giants are poised to accelerate their robotics ambitions, not only deploying robots on factory floors but increasingly manufacturing them at scale. As traditional vehicle sales potentially peak or decline due to demographic shifts, the expertise in mass production and complex assembly will almost certainly be repurposed to build robots. We should expect the bill-of-materials costs to continue their downward trend, meaning security features are increasingly marginalized in favor of market penetration.

The inevitable breach: It is almost certain that we will see a major cyber-physical incident involving embodied AI in the next decade. This could take the form of large-scale operational downtime in a roboticized factory, a legal crisis arising from a hijacked robot causing human injury, or a high-profile case of industrial espionage involving a robot used to map a secret facility. The incident involving Ecovacs vacuums relaying obscenities and racial slurs from a remote hacker is an early indicator of how these risks may evolve.

A new security industry: The next decade will likely see the rise of a dedicated industry focused on securing humanoid robots. Just as the PC era gave birth to antivirus software and the cloud era to SASE, the robotic era will require specialized firms that can provide "physical firewalling," behavioral motor-control monitoring, and "robot-specific" threat intelligence. Companies such as Periphery are examples of where the industry could be rapidly headed.

Mitigation Level	Action Item	Responsibility
Strategic	Conduct a comprehensive risk analysis before deploying embodied AI.	Board of Directors / CISO
Tactical	Segment networks and monitor for outbound telemetry to foreign IPs.	Network Security Team
Operational	Establish playbooks for "emergency robot removal" during a breach.	Operations Manager
Procurement	Verify "security by design" and the absence of hard-coded keys.	Engineering/ Security Audit

Table 2: *Mitigation strategies for business risks, with recommended ownership (Source: Recorded Future)*

Mitigations

Monitor and maintain a vulnerability register: Track disclosed vulnerabilities in any robotic platform your organization deploys or is considering. Establish a playbook for quickly patching or taking robots offline, and understand the operational downtime cost before, not after, a vulnerability is discovered. Recorded Future Vulnerability Intelligence can provide continuous monitoring of emerging CVEs and disclosures specific to embodied AI platforms.

Communicate procurement risks to decision-makers: If your company is purchasing robotics for its operations, the risks of surveillance, covert data exfiltration, and remote compromise must be clearly documented and escalated to the board. Purchasing decisions based solely on unit cost, without accounting for tail risk, are not cost-effective.

Interrogate manufacturers on security by design: Work as closely as possible with manufacturers to understand what security measures are built into the platform, what telemetry is collected, where it is routed, and how firmware updates are managed. If responses are unsatisfactory or opaque, treat that as a material factor in procurement. If the decision-makers' risk appetite remains high regardless, document it formally.

Monitor the macro landscape continuously: New manufacturers are entering the market at a rapid pace, some with security as an afterthought. Recorded Future Threat Intelligence and Geopolitical Intelligence can assist organizations in tracking which companies are emerging, which have ties to state interests, and how the regulatory environment is shifting in key jurisdictions.

Further Reading

SOURCE	TITLE
Insikt Group	The Future of Humanoid Robotics
Sellafield Ltd	Sellafield Technology Radar: Robotics
Alias Robotics	Cybersecurity AI: Humanoid Robots as Attack Vectors
Bin4ry	Unitree Go1: Who is speaking to my dog?
Bank of America	Humanoid robots 101

Risk Scenario

Scenario: ACME Ltd manufactures high-grade munitions for a Western military and allied customers. To reduce human exposure to dangerous materials, it buys 100 humanoid robots from an overseas vendor. The chosen model is already used in comparable factories abroad and costs roughly half as much as an alternative sourced from the United States.

First-order Implications

Threat

The robots arrive and are deployed. Within seconds of powering on, each unit establishes a persistent connection to the telemetry servers in the country of manufacture. Camera feeds, microphone audio, LiDAR spatial data, and GPS coordinates continuously begin routing abroad, providing a granular, real-time map of the facility: its layout, chemical inventory, production processes, and personnel.

Risk

Competitive Disadvantage: Sensitive production methods and munitions designs are transmitted to a state-sponsored threat actor capable of using them both to replicate capabilities and to undercut ACME Ltd on future contracts via state-owned competitors.

Second-order Implications

Threat

A sensitive area of the factory has not been cleared for robot access. Operatives within Bluetooth range transmit a command to a single unit, instructing it to enter the restricted zone. The command propagates to all robots within ten meters. Human operators notice the unexpected movement and raise the alarm. Within hours, the incident reaches the press.

Risk

Brand Impairment: Reports emerge that ACME Ltd has suffered a security incident involving rogue robots moving autonomously through its facility. This is particularly damaging, given that the rationale for deploying robots was explicitly to keep human workers safer.

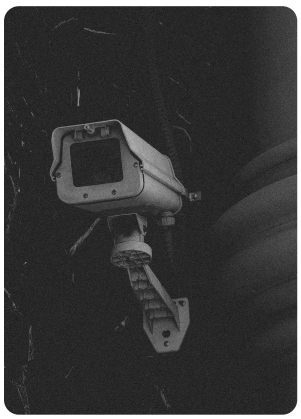

Third-order Implications

Threat

Security experts are brought in to assess the robots. They confirm that every unit has been transmitting camera, microphone, and GPS telemetry to servers in Asia since deployment. They also confirm that the Bluetooth vulnerability has been actively exploited. ACME Ltd immediately pulls all 100 robots from the factory floor.

Risk

Operational Downtime: The factory shuts down while robots are removed and human operators are reinstated at high cost. Output ceases for weeks.
Legal and Compliance: Regulators investigate ACME Ltd for failure to apply adequate security measures. Fines follow.
Contract Loss: The government client cancels contracts with ACME Ltd, citing the national security implications of the breach. The reputational damage extends to the company's broader client base.

Key
Legal or compliance failure: Breach of laws, regulations, or industry standards resulting in liability or sanctions.
Operational disruption: Interruption to normal business processes affecting productivity or service delivery.
Brand impairment: Damage to reputation that reduces customer trust and market value.
Financial fraud: Unauthorized manipulation or theft of financial assets for personal or organizational gain.
Competitive disadvantage: Loss of market position due to inferior capabilities, intelligence, or innovation.

References:
The Risk Business: Second Edition
Intelligence to Risk
The Intelligence Handbook