



www.pipelinepub.com

Volume 22, Issue 9

Rise of the Robots: You Wouldn't Let a Stranger Roam Your Building. Why Let a Robot?

By: [Koroush Saraf](#)

Walk into any modern warehouse, hospital, hotel, airport, campus, or corporate lobby, and you will increasingly find a non-human colleague.

A floor-cleaning robot navigating the atrium.
An autonomous delivery cart moving supplies.
A security drone patrolling the perimeter.
A robotic kiosk greeting visitors.
A connected machine performing work that used to require a person.



We have decided, almost overnight, that autonomous machines belong inside some of our most sensitive physical spaces. What we have not decided is how to vet them.

That gap should concern every enterprise leader, network operator, security team, and risk officer.

Consider how seriously we treat human access. No serious organization lets an unbadged stranger walk freely through its facilities. We run background checks. We issue credentials. We restrict access by role. We log entry and exit. We assume, correctly, that physical presence inside our walls is a privilege that must be earned and continuously verified. Then we wheel in a 200-pound autonomous machine packed with cameras, microphones, sensors, motors, network radios, cloud connectivity, and a software stack most enterprises have never inspected — and we plug it in.

No background check.

No behavioral credentialing.

No independent validation.

No clear view into what it senses, where it sends data, what commands it can receive, or what it does when no one is watching.

This is essentially an unmanaged, third-party operational technology risk. And it is about to become a very expensive one (for whom?).

The threat is not hypothetical, and it comes from three directions at once.

The first is the vendor. A robot you purchased is often still, in a meaningful sense, theirs. It phones home. It receives updates. It sends telemetry. It may depend on a cloud backend, a mobile app, a remote support tunnel, or a fleet-management API that the enterprise does not control. That means the vendor, intentionally or unintentionally, has a path into your facility. In some cases, so do the vendor's subcontractors, cloud providers, support teams, and software supply chain.

The second is the attacker. Many connected machines are deployed like ordinary equipment, but they are not ordinary equipment. They are mobile computers with sensors, actuators, and network access. Some run aging operating systems, open-source robotics middleware, exposed APIs, weak authentication, poorly segmented connectivity, or remote access functions designed for convenience rather than security. If an attacker compromises a fleet-management platform, cloud API, update channel, or command-and-control path, the result is not just a data breach. It can become a coordinated physical intrusion.

The third, and most underrated, is the accident. Autonomous systems fail in ways their makers did not fully test. A navigation model misreads an edge case. A sensor is spoofed. An actuator fires at the wrong time. An AI system interprets a command too broadly. A robot behaves correctly according to its internal logic, but dangerously in the real world. These failures do not require malice to create harm.

That is what makes robots different from laptops, cameras, or access points. A compromised robot can observe, move, block, carry, damage, distract, and interact with the physical world. It is not just another endpoint. It is an endpoint with a body and capabilities.

The robot is the tip of the iceberg

The robot is only the visible version of a broader problem. The same issue is already entering the enterprise without wheels. Autonomous AI agents are being connected to CRM systems, support platforms, code repositories, databases, financial workflows, HR systems, ticketing systems, and operational tools. These agents may not roam the hallway, but they operate inside the enterprise trust boundary. They have credentials. They take actions. They make decisions. They access data. They call APIs. They can trigger workflows.

From a risk perspective, an autonomous AI agent is another actor inside your environment. It has the same three exposures as the robot.

A vendor may be able to reach it.

An attacker may be able to hijack it.

And the system itself may do something no one intended.

The uncomfortable truth is this: a robot, an AI agent, or any autonomous connected system is not simply a product. It is continuously operating third-party operational technology living inside your trust boundary.

And right now, most organizations are not checking these systems with anything close to the rigor they apply to people, networks, applications, or traditional infrastructure.

That has to change.

Real validation needs to examine every layer

At the hardware layer, enterprises need to understand what the machine physically is and what it can do. What sensors does it carry? Does it have cameras, microphones, lidar, radar, GPS, cellular radios, Wi-Fi, Bluetooth, USB ports, or local storage? What can it touch, move, lift, open, block, or damage? Can components be tampered with? Does the hardware include undocumented radios, debug ports, or maintenance interfaces? What is the blast radius if the machine malfunctions or is compromised?

At the software layer, organizations need to inspect the operating system, firmware, middleware, dependencies, update process, authentication model, access controls, logging, and vulnerability history. A robot fleet should not receive less scrutiny than an application server simply because it has wheels. The enterprise needs to know what software is running, who maintains it, how it is patched, how updates are signed, and whether the system can be independently tested.

At the communications layer, enterprises need to know exactly how the system connects and communicates. Is it using enterprise Wi-Fi, private cellular, public cellular, Bluetooth, satellite, or a combination? Where does traffic go? Which cloud services receive telemetry? Are control commands routed through the public internet? Is traffic encrypted end-to-end? Is a private tunnel required? Can the vendor remotely access the device? Can the command-and-control channel be isolated, monitored, restricted, or shut down?

At the command-and-control layer, the enterprise needs independent validation of who can issue commands, how those commands are authenticated, what actions are allowed, and whether command paths can be hijacked or abused. This layer matters because it is where digital compromise can become physical action. A weak command-and-control model can turn a useful robot fleet into a coordinated intrusion platform.

At the AI and behavior layer — the hardest and most neglected layer — organizations need to validate what the system actually does under real-world conditions. How does it behave when sensors are noisy? When instructions conflict? When adversarial inputs are introduced? When a user gives it an ambiguous command? When it reaches an edge case the vendor demo never covered? Does the system stay within approved boundaries, or can it be pushed into unsafe, biased, unauthorized, or operationally damaging behavior?

The industry needs a new discipline: independent IoT command-and-control validation

Traditional security testing is necessary, but it is not enough. Compliance questionnaires are necessary, but they are not enough. Vendor promises are necessary, but they are not enough. Enterprises need a third-party validation service that can test autonomous and connected systems across hardware, software, communications, command-and-control, and AI behavior before they are

admitted into production environments.

This validation should not be a one-time checklist. These systems change continuously. Vendors push updates. AI models evolve. Fleet-management platforms add features. Connectivity paths change. New APIs are exposed. A system that was safe when deployed can become risky after a software update, cloud change, configuration drift, or integration with another enterprise system.

That is why command-and-control validation has to be continuous.

Before deployment, validate the system.

During operation, monitor the system.

After updates, re-check the system.

When behavior changes, investigate the system.

When access expands, re-credential the system.

Telecom and connectivity industries have an important role to play

As robots, autonomous agents, and connected machines scale, they will depend on reliable, secure connectivity: private 5G, enterprise wireless, cellular IoT, edge compute, secure tunnels, device identity, policy enforcement, and managed connectivity. The network is no longer just a transport layer. It becomes part of the control plane for trust.

That creates both responsibility and opportunity.

Operators and connectivity providers should not only connect these systems, but they should also help enterprises control, segment, monitor, and validate them. The next phase of enterprise connectivity will not be defined only by speed, coverage, or cost per bit. It will be defined by whether the network can help prove that every connected actor is known, authorized, constrained, observable, and safe.

That applies to robots in buildings.

It applies to autonomous vehicles in yards.

It applies to drones on campuses.

It applies to AI agents in enterprise systems.

It applies to every connected machine that can sense, decide, communicate, and act.

More robots are coming

The future will not be less autonomous. It will be more autonomous. More robots. More agents. More connected machines. More third-party systems operating inside enterprise environments. More systems that bridge the digital world and the physical world.

So the principle has to be simple:

Every robot, every agent, and every autonomous connected system admitted into your environment is a new actor inside your trust boundary. Treat it accordingly.

Credential it.

Constrain it.

Segment it.

Validate its hardware.
Validate its software.
Validate its communications.
Validate its command-and-control path.
Validate its AI behavior.
Monitor it continuously.
Re-check it when it changes.

We learned long ago not to trust an unverified stranger inside the building.

The robots and agents deserve no more benefit of the doubt.

Not for distribution or reproduction.