ALLIANCE FOR INNOVATION AND INFRASTRUCTURE

Analysis

The Robot Is Not the Problem

Physical AI and the value of data architecture in critical infrastructure

May 2026

Independent, nonpartisan research to inform policy and drive progress.

Mothusi Pahl

Aii.org

The Robot Is Not the Problem

Physical AI and the value of data architecture in critical infrastructure

The Devices and the Data

Autonomous robots are being deployed across U.S. infrastructure at scale. Pipeline inspection crawlers run inside crude lines. Substation patrol robots run scheduled inspection rounds at closed industrial sites. Autonomous haul trucks cross private mining roads. Bridge inspection drones map structural details across state Department of Transportation (DOT) programs. Container straddles handle cargo moves at the largest U.S. ports.[1] Each emerging robotic application runs under a contract that names what the technology will deliver.

These deployments raise the standard concerns: data exfiltration, surveillance backdoors, and remote hijacking. These concerns are legitimate. But giving priority to the question "who manufactured the robot" is evidence of thinking in terms of a 2010s threat model. The pressing questions today are about the data collected across all phases of operations. In aggregate, these expanding deployments will produce a near real-time perceptual record of U.S. critical infrastructure. This data collection is a function of operation, regardless of intent. Every physical AI that performs a task perceives its environment. And every perception produces a record, regardless of the task in the contract.

Throughout this piece, two terms carry specific meanings. "Operator" means the infrastructure owner that procures and deploys the robot (the port authority, the airport, the state DOT, the utility, the pipeline company, the refinery). Where a separate contractor physically runs the asset, the operator is still the owner who signed the deployment contract and bears the consequences of its terms. "Architecture" means the relationship between what the robot perceives, where that perception data flows, and what records result. The argument of this piece is that the architecture, and the sensitivities baked into it, are determined by what the robot does, not by who built it.

In December 2025, the FCC issued a National Security Determination blocking new authorizations for foreign-made aerial drones across categories including surveillance, transmission-line inspection, pipeline monitoring, post-storm survey, etc.[2] The FCC's reasoning rested on supply-chain and data-transmission risks that apply to every foreign-made unmanned aircraft, whatever it is doing. The architectural reasoning and risks for autonomous robots are similar to drones but run much deeper. In contrast to human-controlled teleoperated devices, perception data in autonomous systems is mandatory, continuous, and independent of the

mission. This makes physical AI operations and intelligence collection hard to tell apart, especially when foreign or foreign-influenced operating systems are ingesting the perception data.

The Architecture Produces the Exposures

Robots in surveillance and security are the application categories that will draw the most OEM and AI-company commitment, not because the deployments are uniquely valuable but because they generate data at scale, and because the data are collected on private property, outside of most consumer-data and public-space regimes. This is where the data architecture shows up first. The permission structure and operating environment make data capture in surveillance and security applications easier to scale than in any other category. The AI industry calls this "the data flywheel": More deployed robots produce more training data, better models, and broader deployment.[3]

These surveillance and security applications are also where political attention has landed first. The same data collection conditions and sensitivities are present at the pipeline company, the port authority, the state DOT, the airport authority, and the refinery. All are operators under the definition above, and all are deploying autonomous systems into critical infrastructure under the same architectural conditions and on timelines that are part of the coming surveillance and security wave.

But the records that feed these autonomous robot training pipelines are not unique to surveillance. Every deployed system generates an operations record that the manufacturer wants to aggregate into the same architecture, because this accumulated record is what trains the OEM's next model.[4]

Within this architecture there are two streams of data to account for. The first data stream connects directly to the product or the service: the inspection records, the haulage logs, the security feeds, the inventory counts. These are the actual deliverables named in the contract. The second data stream is the operations record: the environmental maps, the personnel detection, the route memory, the acoustic signatures, the Bluetooth and WiFi device signatures the sensors pick up, and more. Together these elements include everything the system perceived in order to perform its deliverable, plus everything else its sensors picked up along the way. Robot OEMs then do three things with these two data streams.

The first output is the deliverable. These items are named in the contract, evaluated in procurement, tracked on dashboards and accounted for in the operations budget. A pipeline inspection crawler delivers a wall-thickness map and a defect log. An autonomous haul truck delivers tonnage moved. A substation patrol robot delivers thermal anomaly reports and intrusion

alerts. An autonomous baggage tug delivers turnaround timing and bag-cart routings. A container straddle delivers cargo move counts and yard positions. These are the data the operator specified. The procurement officer has a clear view of it. The contract governs it and it is the part of the deployment that the operator fully understands and expects.

The second output category is training data. Every hour the deployed fleet operates, the robot OEM takes in what its systems perceive. A pipeline inspection crawler in West Texas improves the in-pipe locomotion and weld-recognition models for crawlers deployed in Alberta. The operator paid for the work delivered at their own facility, but the improvements are deployed across every site the OEM serves.[5]

The third output category sits at the manufacturer level. This is what the deployed fleet has observed and aggregated from every operating environment it ran through. Where the second output improved the manufacturer's models, this third record is the asset itself: equipment configurations across an industry, operational tempo across a sector, infrastructure layouts across a country. Acoustic baselines distinguish specific compressor models; RF signatures specify control-system protocol families; thermal footprints indicate transformer load profiles. The human-behavioral subset of this category ranges from simple foot-traffic and facial recognition to heart rate and blood pressure metrics, all currently sold as commercial products.[6]

Boston Dynamics now markets fleet-management software (Orbit) that aggregates data from all of its deployment sites into centralized dashboards. The company also markets a consulting practice that draws on its test fleet and deployed robots to build bespoke machine-learning systems for customers. These separate offerings show the viability and value of turning these companies' accumulated records into services.[7]

Operators consent to this. Many service agreements grant the manufacturer training rights as a condition of deployment.[8] Every facility that hosts a robot contributes to collective AI training that the OEM owns. As a result, most of the third output category today sits as an unpriced asset on the OEM's books, and the operator has no default right to participate in the decisions about whether and how to productize this category of data.

The operator has no natural visibility into where the training data runs: what cloud infrastructure processes the data, in what jurisdiction, and under whose legal authority. The big three U.S. hyperscalers (AWS, Microsoft, Google) hold roughly two-thirds of the global cloud infrastructure services market.[9] U.S. legal authorities can reach this data through subpoena, through the CLOUD Act (which reaches data held by U.S. providers regardless of where the compute sits[10]), and through National Security Letters (which require no judicial review)[11]. Other states have their own instruments.

Even an OEM operating in good faith does not control the full stack (the cloud service, the orchestration layer, the firmware, the silicon) beneath its product. The data collected by physical AI often moves through layers the contracts do not name and the vendors do not warrant. The same collected sensor data flows into the deliverables, the training data, and the manufacturer's accumulated record, making the three outputs inseparable by default.

On-device-only processing is technically possible. NVIDIA's Jetson is the standard reference platform.[12] But none of the major commercial OEMs currently offer it as a default.[13] Every layer in the stack (from the OEM's data handling down to the silicon) can be addressed contractually and every layer has a price. Operators face an economic question about which exposures they will pay to minimize and which they will accept because the cost of mitigation exceeds the cost of the exposure.

A parallel market for these data also operates without subpoena. In 2024, the New York Times reported that General Motors had been transmitting per-trip driving behavior (hard braking, late-night driving, speeds above 80 mph) to data brokers LexisNexis and Verisk, who sold the data to insurance companies who used it to set rates. Drivers only learned they had been tracked when their premiums rose or they ordered a LexisNexis report.[14] The same pattern will carry robotic infrastructure data to insurers, competitors, activist groups, and law firms - without subpoena and without notice. The operator controls neither the legal pathway nor the commercial one.

Three exposures flow from this record.

> **National security exposure to the country**: the layout and operational tempo of every facility the fleet observed, in one searchable index.
>
> **Commercial exposure to the operator**: throughput patterns and process intelligence that competitors would pay for.
>
> **Personal privacy exposure to the people the system perceived**: the face, gait, and movement record of every worker, contractor, and member of the public the system perceived.

None of these depend on the system being deployed for surveillance. All three flow from the system simply being deployed and operating.

The Ratchet

This pattern is not new. It recurs across the history of surveillance technology. The pattern has a consistent outcome: major commercial surveillance capabilities deployed at scale for a narrow, defensible purpose often expand well beyond their original scope.

Two cases illustrate the pattern.

Amazon's Ring doorbell camera was originally deployed as a home security product. Amazon later launched the Neighbors app, aggregating footage from millions of Ring devices into a neighborhood surveillance network. Law enforcement agencies in more than 2,000 jurisdictions partnered with Ring to request footage through user-targeted requests rather than through warrant processes.[15] The partnership with law enforcement is not the issue. The issue is that a Ring camera records everyone within view (visitors, houseguests, children), none of whom consented to being recorded by Amazon or to having that footage sit on Amazon's servers. The data sits available for Amazon's business purposes and is subject to any legal instrument that compels third-party data disclosure.

Cell phone location data followed a similar trajectory. Wireless carriers collected location data for network management and billing. Data aggregators purchased this data from carriers and resold it to commercial buyers, including bail bond companies and bounty hunters. In 2018, a sitting Missouri sheriff was found guilty of using such a service to track a judge and fellow law enforcement officers.[16] No cell phone client consented to these downstream uses when they signed their original wireless contract. The data was collected for one original purpose, then monetized in dozens of others and regulation has not caught up.

Deployed robots will follow the same trajectory, faster and across the entire commercial robotics market. The data is richest in the security and surveillance segments where a continuous record of human movement, equipment-level detail and three-dimensional spatial coverage captures a massive amount of detail and variability. Every robot produces some form of this data structurally. A warehouse humanoid, a delivery vehicle, an inspection crawler and a patrol quadruped feed the same data architecture. The ratchet does not reverse. The question is how far it advances before governance catches up.

Exposures and Contemplated Legislation

The American Security Robotics Act, introduced in March 2026 by Senators Cotton and Schumer, is the first federal bill specifically addressing unmanned ground vehicle systems in terms of national security.[17] It is a sound piece of industrial policy and a credible response to one specific risk. Unfortunately, the bill is also a repeat of a recurring template: Government

response arriving roughly one technology generation behind operational reality, and further behind the frontier that the commercial market will reach in the next 12–18 months.

The Cotton/Schumer bill prohibits federal procurement and operation of unmanned ground vehicles manufactured or assembled by entities controlled by China, Russia, North Korea or Iran (using "unmanned ground vehicle systems" as the regulatory term for autonomous ground robots). This is one subset of one risk (manufacturing) and is presented in the form of an individual category of physical devices (UGVs). The bill is silent on the larger fact: Most compute that powers physical AI today is delivered as a service and synchronized to the cloud because the AI models that run these robots are usually too large to fit onboard.

The national security framing also restricts foreign-built autonomous systems while permitting American-built ones to operate with the exact same data architecture. An American flag on the autonomous ground vehicle gets read as a guarantee about the data security of everything downstream of the device. This creates a projection of false security. A U.S.-manufactured chassis running a foreign-trained or foreign-hosted program is, from a perception-data exfiltration standpoint, indistinguishable from a foreign-manufactured chassis running a domestic policy. The bill bans the chassis. It does not address the autonomy stack: where the perception data actually flows.

Beyond what the bill misses diagnostically, there is the ever-present strategic risk of losing access to superior technology. Foreign robotics systems across the stack (hardware, operating systems, software, AI models) are in many cases superior to current American alternatives, leading on price and performance where some American offerings do not compete. Banning these imports on point of origin is a choice with real costs to U.S. operators and to U.S. national security.

The bill's industrial policy and supply chain objectives are sound, but another major gap is the fact that the bill does not speak to the obvious issue of foreign-trained models embedded in U.S. platforms. If Unitree is banned in the U.S. but Unitree’s model weights and training data and policy infrastructure are acceptable for licensing by a U.S. OEM who then integrates those weights into a domestically manufactured chassis, then the bill as drafted is a false positive for U.S. national security.

In earlier work I proposed SAAFE-7[18], a seven-dimension framework for evaluating compute exposure: power, data security, air-gap, physical isolation, supply chain integrity, jurisdiction alignment, and trust. The Cotton/Schumer bill addresses one face of one dimension - manufacturer origin, within supply chain integrity. The bill is not a partial solution to a seven-dimension problem. It is a solution to one face of one dimension.

What to Do Now

What follows assumes the contract leverage of a private operator (utility, port, refinery, mine). Federally funded infrastructure faces the same exposures but reaches them through 8(a) and OTA pathways that constrain procurement authority; the prescriptions translate with friction.

This is not a security decision. It is a decision about who owns the recording when the contract ends. It applies to every autonomous system you deploy. Every autonomous system generates a continuous record of the facility as a byproduct of operating within the facility. The contract on your desk decides who owns that record, where it is held, who can access it and what the vendor is permitted to do with the record. A contract that does not name these terms surrenders them by default.

What to do. Three actions, in order:

1. Audit the existing footprint within 90 days. Before negotiating the next contract, inventory the autonomous systems already deployed across your sites and pull the contracts that govern them. Most will be silent on the six terms below. Assign a single accountable owner (deputy GC or CISO with a procurement counterpart) and report findings to the board's risk committee. You cannot negotiate the next contract intelligently without knowing what you have already conceded in the prior ones.
2. Make two contractual provisions non-negotiable on every renewal and new agreement. These are the two that, if won, make the others enforceable:
    a. **Operations record ownership.** You own everything the system observes in the course of operating at your facility, not only the deliverables you contracted for. The contract names the record explicitly and assigns title to the operator.
    b. **Training rights.** The vendor may not use your facility's operations record to improve models deployed at other facilities (including peers and/or competitors) without an affirmative, scoped, time-limited license that you can revoke. Default is no.
3. Require four additional provisions. A vendor that agrees to (1) and (2) but refuses these is telling you something.
    a. **Compute and storage location.** Where data is processed, where it is held, which jurisdictions the pipeline touches at every layer.
    b. **Manufacturer's accumulated record.** What the vendor may do with aggregated fleet data and what notice you are owed if that use changes.
    c. **Retention and deletion.** Storage duration, deletion protocols, and what survives in derived models after the contract ends.

d. **Legal-compulsion access.** The conditions under which the vendor will share your data with government agencies — domestic or foreign — and the notice you are owed.

A test for vendor architecture, not just vendor contracts. Contract terms are only as good as the architecture behind them. A vendor that ships data sovereignty as a default (on-device-only processing as a real configuration option, training opt-outs the operator can actually exercise, jurisdictional disclosure as a standard datasheet field) is built for the regulatory environment that is arriving. A vendor that treats this as custom engineering is built for the one that is leaving. Run the architecture question through your CIO in parallel with procurement, not after.

Governance. This is not procurement's problem to solve. The operations record is generated by operations, governed by IT and security, exposed by legal, and material to the board's risk committee. Set a materiality threshold. Above it, every autonomous-system contract goes to a quarterly review with all four functions in the room. The first meeting reviews the audit from action 1.

A vendor that will not agree to provisions (1) and (2) in the master agreement (not a side letter) is telling you which of the three outputs they are actually selling and which of those they expect to keep.

What This Means for OEM Product Strategy

The provisions above are about to become procurement standard. The operators that adopt them (and they will, on a calendar set by the Cotton-Schumer bill, the first state-level Physical AI Surveillance Act, and the first adverse court ruling on platform-collected data) will evaluate OEMs on different terms than they did last year. The dimension that will matter is which of the three outputs an OEM's product architecture treats as the deliverable, which as a byproduct, and which as a strategic asset the operator was never told existed. Every OEM selling into enterprise procurement over the next 24 months (Tesla, Figure, Apptronik, Agility, Boston Dynamics, and the tier behind them)[19] is answering this question in the master service agreement template that goes out with the next deal, whether deliberately or not.

The strategic choice is whether data sovereignty is the default configuration or an opt-in feature. A default configuration means the OEM has built data sovereignty into the standard product: on-device-only processing as a real option, training opt-outs the operator can actually exercise, jurisdictional disclosure across the stack, explicit handling of the manufacturer's accumulated record. Opt-in means the OEM is hedging: selling these capabilities to enterprise customers who

ask for them, preserving the data flywheel for everyone else. The hedge looks rational this year and dangerous in three, because procurement standards do not reverse. The OEM that offers these capabilities at competitive prices wins. The one that prices them as a premium upgrade loses.

The honest tension underneath is that the data flywheel is the moat. The operations record an OEM's fleet generates is what trains the next generation of model and the OEM that constrains that record at the operator's request is constraining its own product roadmap. There is only the choice of which side of that tension the architecture is built around, and that choice is hard to reverse.

The OEM that has shipped a defensible answer in the default configuration has a story to tell in enterprise procurement. The OEM that has not is going to hear the question for the first time from a procurement officer who has read a piece like this one. The contract is the governance decision. The operator who does not press for terms gets the OEM's.

Mothusi Pahl is principal at H&L Energy, where he focuses on infrastructure risk, energy policy, and the intersection of physical systems and compute sovereignty. He developed the SAAFE-7 framework for evaluating sovereignty exposure in cloud and AI training environments, proposed as an open standard with attribution.

Published with the Alliance for Innovation and Infrastructure (Aii). The views and opinions expressed are solely those of the author and do not necessarily reflect the views of Aii or its leadership.

Citations and Notes

1. On the deployment landscape: ANYbotics' ANYmal X has been deployed by PETRONAS and Schlumberger (now SLB) for autonomous oil-and-gas inspection, https://www.anybotics.com/customers/. Boston Dynamics Spot has been adopted by utilities including Avangrid and National Grid for substation inspection, https://bostondynamics.com/case-studies/. Caterpillar's autonomous haul truck program reported 690 deployed Cat 793F autonomous units globally as of end of 2024, with a target of 1,300 units by 2030, https://www.caterpillar.com/en/news/corporate-press-releases/h/2024/cat-autonomous-mining-trucks-milestone.html. Minnesota DOT's drone fleet had grown to 33 UAS by 2025 across the agency's bridge inspection program, https://www.dot.state.mn.us/aero/uas/index.html. Long Beach Container Terminal (LBCT) and TraPac at the Port of Los Angeles operate fully automated container handling using straddle carriers and automated stacking cranes, https://www.lbct.com and https://trapac.com.

2. FCC National Security Determination, December 22, 2025, DA 25-1086. The Determination blocked new equipment authorizations for foreign-made unmanned aircraft systems from covered manufacturers. https://www.fcc.gov/document/fcc-national-security-determination-foreign-uas-2025

3. On the data flywheel as the strategic logic of physical AI commercialization: McKinsey, "Humanoid Robots: Crossing the Chasm from Concept to Commercial Reality," October 2025, identifies "embodied data" and the flywheel (better models → better performance → more data) as primary strategic questions for humanoid commercialization, https://www.mckinsey.com. Roland Berger, "Humanoid Robots 2026: The Convergence Moment," April 2026: "The bottleneck in humanoid robotics is training data. Manufacturers trading data access receive cutting-edge technology in return," https://www.rolandberger.com. MIT Technology Review, "The Robot Race Is Fueling a Fight for Training Data," April 2024, https://www.technologyreview.com.

4. On the data flywheel in OEM commercial framing and academic literature: Knightscope describes its products as "Autonomous Data Robots" generating ~90 TB/year per unit (William Santana Li, CEO, InvestorBrandNetwork). Google DeepMind / Agile Robots partnership announcement, March 26, 2026 (The Robot Report, https://www.therobotreport.com). Grannen et al., "Robot-Powered Data Flywheels," arXiv:2511.19647, November 2025, https://arxiv.org/abs/2511.19647

5. Bain & Company, "Humanoid Robots: From Demos to Deployment," 2025, https://www.bain.com/insights/humanoid-robots-from-demos-to-deployment-technology-report-2025/; IEEE Spectrum, "Humanoid Robots: The State of the Art," October 2025, https://spectrum.ieee.org

6. Placer.ai, RetailNext, Sensormatic (Johnson Controls). Grand View Research, "People Counting System Market," 2024. https://www.grandviewresearch.com/industry-analysis/people-counting-system-market-report

7. Boston Dynamics Orbit fleet management software — product page: enterprise users have access to "centralized dashboards that aggregate data from all sites — giving you a unified view of robot activity, site performance, and fleet health for Spot, Stretch, and eventually Atlas." https://bostondynamics.com/products/orbit/. Boston Dynamics Consulting, launched March 2025: the company explicitly draws on "rich data from our test fleet and deployed robots to identify gaps, build robust machine learning pipelines, and deliver practical results with AI for our customers." https://bostondynamics.com/news/boston-dynamics-launches-consulting-services/; https://bostondynamics.com/faq/

8. On training-rights grants in commercial AI/robotics service agreements: Morgan Lewis, "Negotiating AI Provisions in Commercial and Technology Contracts: Where the Market Is Heading," April 2026, identifies IP allocation — including "what rights the customer retains in training data it provides, and whether the vendor can use customer data to improve its model" — as among the most actively negotiated areas. https://www.morganlewis.com/blogs/sourcingatmorganlewis/2026/04/negotiating-ai-provisions-in-commercial-

and-technology-contracts-where-the-market-is-heading. See also "Understanding Training Data in Contracts with AI Vendors," ContractNerds, September 2025, on the prevalence of vendor-friendly default training-data terms in trial agreements, master agreements, and DPAs. https://contractnerds.com/understanding-training-data-in-contracts-with-ai-vendors/

9. Synergy Research Group, "Cloud Market Share Trends — Big Three Together Hold 63% While Oracle and the Neoclouds Inch Higher," November 19, 2025: AWS, Microsoft Azure, and Google Cloud together accounted for 63% of enterprise spending on cloud infrastructure services in Q3 2025, with worldwide market value at $107 billion in Q3. Q4 2025 estimates from Synergy and CRN put the combined share at approximately 68%. https://www.srgresearch.com/articles/cloud-market-share-trends-big-three-together-hold-63-while-oracle-and-the-neoclouds-inch-higher

10. CLOUD Act, 18 USC 2713 (March 23, 2018). https://www.congress.gov/bill/115th-congress/house-bill/4943/text

11. FISA Section 702, 50 USC 1881a. https://www.law.cornell.edu/uscode/text/50/1881a

12. Xu et al., "Empowering Edge Intelligence: A Comprehensive Survey on On-Device AI Models," ACM Computing Surveys, 2025: NVIDIA's Jetson platform is identified as the canonical hardware platform for local inference in industrial robots, autonomous vehicles, and IoT systems, supporting "real-time data processing" without cloud dependency. https://dl.acm.org/doi/full/10.1145/3724420; arXiv preprint at https://arxiv.org/pdf/2503.06027. NVIDIA's own Jetson / Isaac Lab / GR00T documentation describes sub-30 ms transformer-based inference on the edge as the standard pattern for modern robotics deployment. https://developer.nvidia.com/blog/getting-started-with-edge-ai-on-nvidia-jetson-llms-vlms-and-foundation-models-for-robotics/

13. Author's review of OEM product documentation, public marketing materials, and service agreement summaries (April 2026), covering: Knightscope, Boston Dynamics, ANYbotics, Cobalt Robotics, Tesla (Optimus), Figure, Apptronik, Unitree, AgiBot, UBTECH. Methodology: review of public product pages, technical documentation, marketing collateral, and where available service agreement templates and customer testimonials. None of the reviewed OEMs offers on-device-only processing as the standard (default) product configuration; cloud-dependent processing and manufacturer training rights are uniformly the default.

14. Kashmir Hill, "Automakers Are Sharing Consumers' Driving Behavior With Insurance Companies," New York Times, March 11, 2024. The article documented that General Motors transmitted per-trip driving data — hard braking, late-night driving, speeds over 80 mph — from OnStar Smart Driver to data brokers LexisNexis and Verisk, which sold the data to insurance companies (eight companies had requested driver Kenn Dahl's data within one month). https://www.nytimes.com/2024/03/11/technology/carmakers-driver-tracking-insurance.html. Senators Wyden and Markey requested an FTC investigation in August 2024; GM ended the OnStar Smart Driver program in April 2024 and terminated its data-sharing partnerships with LexisNexis and Verisk in March 2024. https://www.repairerdrivennews.com/2024/08/01/two-more-congress-members-ask-ftc-to-investigate-alleged-consumer-data-selling-by-automakers/

15. Ring letter to Sen. Ed Markey, 2022 (>2,000 partnerships), https://www.markey.senate.gov/news/press-releases/senator-markeys-probe-into-amazon-rings-surveillance-practices-and-cooperation-with-police; LA Business Journal, February 2024 (~2,680 law enforcement, 622 fire by 2024), https://labusinessjournal.com/featured/ring/

16. Jennifer Valentino-DeVries, "Service Meant to Monitor Inmates' Calls Could Track You, Too," New York Times, May 10, 2018. https://www.nytimes.com/2018/05/10/technology/cellphone-tracking-securus.html.

ACLU, "Company That Handles Prison Phone Calls Is Surveilling People Who Aren't in Prison," https://www.aclu.org. Former Mississippi County, Missouri Sheriff Cory Hutcheson.

17. Senators Tom Cotton (R-AR) and Chuck Schumer (D-NY) introduced the American Security Robotics Act in March 2026. The bill targets unmanned ground vehicle systems manufactured or assembled by entities subject to control by China, Russia, North Korea or Iran. https://www.govinfo.gov/app/details/BILLS-119s4235is

18. Pahl, Mothusi, "The Sovereignty Gap in Compute," H&L Energy, 2026. SAAFE-7 (Sovereignty Architecture Assessment Framework for Exposure) covers seven dimensions: power, data security, air-gapped, physically isolated, supply chain integrity, jurisdiction-aligned, trust. Open standard with attribution. Supply chain integrity is the dimension least directly addressed by the checklist below. The body's discussion of "the cloud service, the orchestration layer, the firmware, the silicon" is where that work lives in the current article; a fuller treatment of the four named dimensions and their cross-cutting application to deployed robotics is the subject of follow-up work.

19. Tesla Q4 2025 earnings call, January 28, 2026, https://www.fool.com/earnings/call-transcripts/2026/01/28/tesla-tsla-q4-2025-earnings-call-transcript/; Electrek, "Musk admits no Optimus robots are doing 'useful work' at Tesla," January 28, 2026, https://electrek.co/2026/01/28/musk-admits-no-optimus-robots-are-doing-useful-work-at-tesla-after-claiming-otherwise/. Knightscope (NASDAQ: KSCP), continued commercial expansion. Unitree IPO filing, China STAR Market, March 20, 2026, $608M raise targeting 20,000 humanoid units in 2026: https://humanoidapac.ai/article-unitree-ipo-2026-en.html; https://www.caixinglobal.com/2026-03-21/unitree-robotics-files-for-608-million-star-market-ipo-102425491.html