

Race to Infinite Labor: The Humanoid Hypothesis

Equity Research
Industrials | Energy and Power Tech-
nologies | Technology

May 26, 2026
Industry Report


[Ross Sparenblek](#) +1 312 364 8361
rsparenblek@williamblair.com

[Jed Dorsheimer](#) +1 617 235 7555
jeddorsheimer@williamblair.com

[Sebastien Naji](#) +1 212 245 6508
snaji@williamblair.com

[Mark Shooter](#) +1 617 235 7593
mshooter@williamblair.com

[Sam Karlov](#) +1 312 551 4586
skarlov@williamblair.com



Please refer to important disclosures on pages 41 and 42. Analyst certification is on page 41. William Blair or an affiliate does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. This report is not intended to provide personal investment advice. The opinions and recommendations herein do not take into account individual client circumstances, objectives, or needs and are not intended as recommendations of particular securities, financial instruments, or strategies to particular clients. The recipient of this report must make its own independent decisions regarding any securities or financial instruments mentioned herein.

Contents

Humanoids, Want or Need?	3
The New Arms Race: National Security Concerns De-risk Path to Commercialization	7
Major Bottlenecks to Scaling Humanoid Robotics	14
From Prototypes to Millions: Milestones and KPIs to Monitor	19
Framing the Opportunity	21
Ways to Play Humanoid Robotics	22
Layers of Humanoid Technology Stack	26
Leading Humanoid OEMs	30
Appendix: Tesla’s Optimus	32

Humanoids, Want or Need?

Industrial productivity increases in discrete steps, each defined by a technology that expands the amount of work one person can do. The loom, steam engine, electricity, production line, and industrial robotics all increased the leverage on human labor. We argue that we have hit a lull in revolutionary leaps in manufacturing technology; instead, we have relied on shifting work globally to where human labor is the cheapest. Western countries then focused on energy-light sectors like software, providing an illusion that manufacturing has become highly efficient and abundant, but in reality, the mines and factories have moved out of their backyards—and out of sight, out of mind.

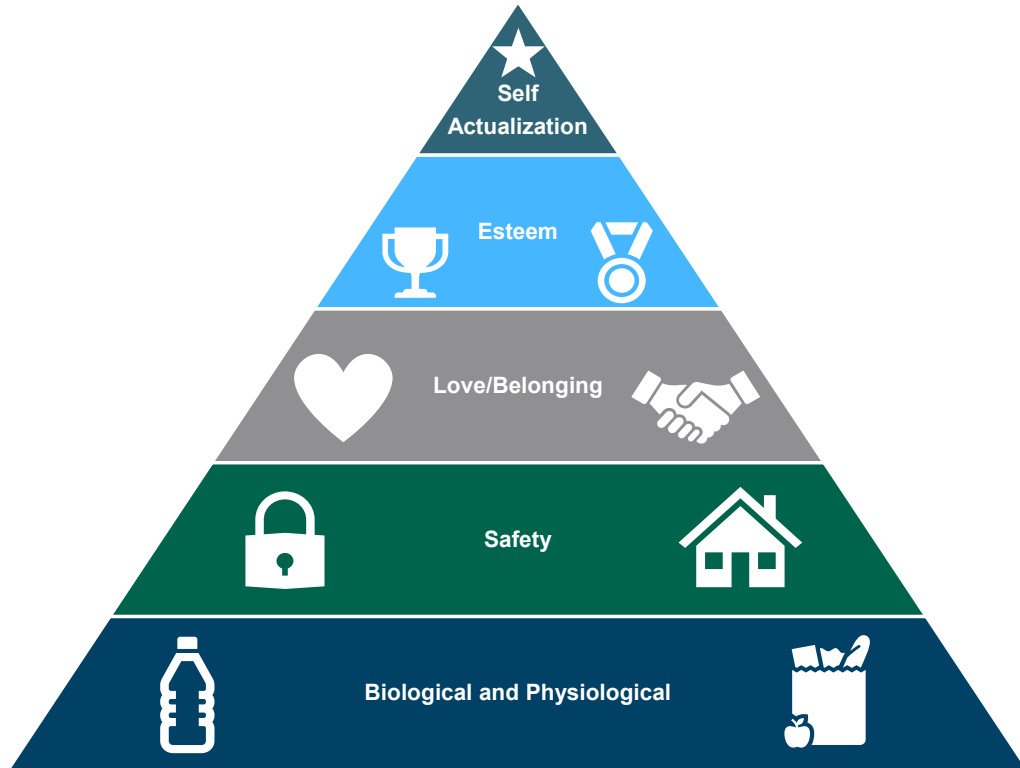
Artificial intelligence has already reshaped digital work. What remains largely untouched is the physical layer of the economy. Goods still need to be moved, assembled, repaired, and maintained. Infrastructure still needs to be built and serviced, warehouses still need picking, and factories still need tending. AI models can plan and reason, but they do not execute in the physical world. The next productivity step requires closing that gap.

Humanoid robotics offer the next revolution in productivity, eliminating the labor cost component of the equation to allow capital- and labor-heavy industries to be redistributed globally again. Traditional industrial robotics are evolutionary, are designed task-specific, and require workflows to be redesigned around them. Humanoids are biomimetic, with the goal to be task-agnostic and drop-shipped into the existing world around them. If successful, the addressable market is not a single function or vertical, but the entire surface area of physical labor.

The Want: Human Labor Liberty

The natural appeal of humanoids is to unburden humans of unfulfilling labor and enable us to devote our finite time to more meaningful pursuits. Viewed through Maslow's hierarchy of needs, removing routine physical labor allows time and energy to shift toward higher-value activities like sciences, arts, athletics, and leisure. A meaningful portion of human time is spent on repetitive, low-satisfaction physical tasks. Entire sectors of the economy rely on mundane and unfulfilling tasks, like cleaning, carrying, sorting, stocking, maintenance, assembly line manufacturing, as well as routine household chores like laundry, dishes, and vacuuming. The value proposition is not just cost reduction, it is time recovery and reallocation to maximize life fulfillment.

Exhibit 1
Race to Infinite Labor: The Humanoid Hypothesis
Maslow's Hierarchy of Needs




Sources: William Blair Equity Research

The Need: Improved Productivity to Meet Structural Concerns

The stronger argument is not that humanoids are desirable, but that they are necessary to increase the productive function to grow the U.S. economy fast enough to sustain our national debt and prevent a debt spiral and runaway inflation. At the same time, we are entering a global decoupling as the rise of China contends with U.S. hegemony, creating a bifurcated and multipolar geopolitical environment. This reshoring, or dual sourcing of supply chains, further strains productivity, as we trade back efficiency for redundancy to rebuild these energy-intensive industries domestically. Unfortunately, when we need it the most, the labor supply is tightening, populations are aging, and workforce participation is flattening or declining.


Exhibit 2
Race to Infinite Labor: The Humanoid Hypothesis
Total U.S. Federal Debt



Sources: U.S. Department of the Treasury

OECD data show the old-age dependency ratio rising from about 19% in 1980 to approximately 31% today, with projections approaching 50% by 2060. The U.S. is faring slightly better than Europe and Japan due to migration, but the trend is clear for all Western countries. The U.S. Census Bureau indicates that by 2030 all baby boomers will be over 65, and by 2034 older adults are expected to outnumber children for the first time.

Exhibit 3
Race to Infinite Labor: The Humanoid Hypothesis
Old Age Dependency Ratio



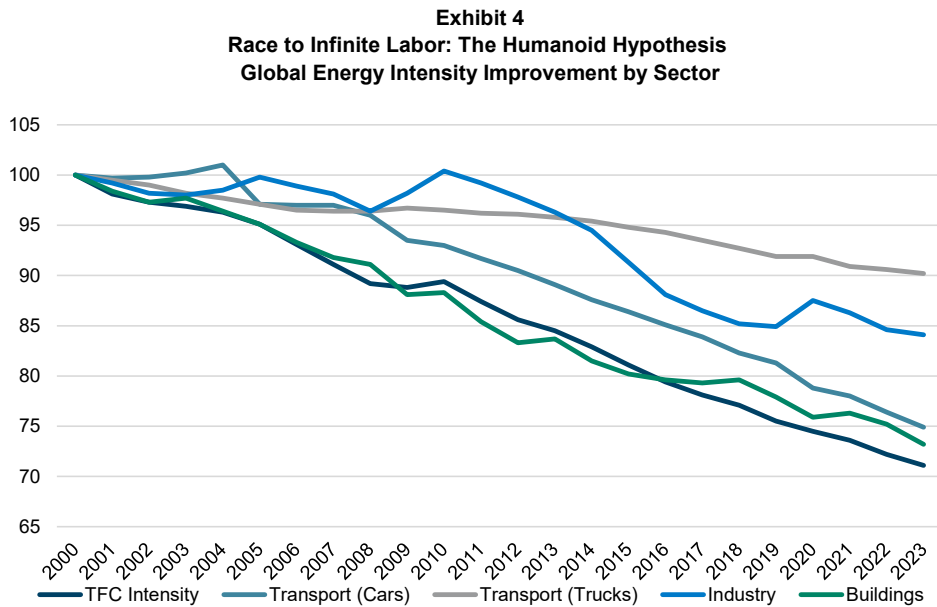
Note:

1) Represents population over the age of 65 as a percentage of the percentage of population between 15 and 64 years of age

Sources: United Nations, Department of Economic and Social Affairs, Population Division (2026)

The challenges in China are even more acute, as decades of the one-child policy have put the country on course to see the largest absolute population decline of any country through 2050. While the two-child policy was implemented in early 2016, the four-year COVID lockdown has led to an average fertility rate of just 1.08 over the last 5 years (2020 to 2024) and just 1.01 in 2024 (UN Population Division, World Population Prospects). This is a structural issue, not cyclical, and requires durable demographic changes across decades to influence.

At a high level, economic output can be simplified into four inputs: energy, labor, capital, and productivity. To the first point, technological advancements have steadily reduced energy intensity globally over the 21st century, as shown in exhibit 4. However, in addition to the shift toward redundant manufacturing, demand for data centers is beginning to put immense pressure on the grid. Energy can only grow as fast as the slowest component of the supply chain, with grid interconnects, transformers, and switch gears now experiencing multiyear lead times. In our view, this leaves productivity as the part of the equation with the most torque.



Note:

1) indexed to the year 2000

Sources: IEA, Global energy intensity improvement by sector in the Net Zero Scenario, 2000-2030

Humanoids may not be a want after all, but a need for the United States to dramatically increase the productive function of our economy. In a fracturing geopolitical landscape, humanoids may be necessary to help rebuild domestic, asset-heavy and energy-intensive industries. This structural necessity is also rapidly intersecting with geopolitical urgency, transforming humanoids from a productivity tool into a strategic national asset.

As the United States seeks to reshore critical manufacturing and restore economic resilience, it finds itself in an intensifying global competition with China over who will control the “infinite labor” platform of the 21st century. What began as a response to domestic labor shortages and debt dynamics has evolved into a full-spectrum robotics arms race, where technological leadership in embodied AI directly translates into economic sovereignty, supply-chain security, and military advantage.

The New Arms Race: National Security Concerns De-risk Path to Commercialization

Private sector investments in robotics and embodied AI have skyrocketed over the past 24 months, led by humanoid robotics, as breakthroughs in large language and world models have dramatically accelerated the timeline to mass commercialization. This has fueled market optimism, as evident in a growing number of product demonstration videos and companies outlining their viable economics across the supply chain. Skeptics, however, point to the technologies' nascent stage, limited real-world deployments, and often lofty valuations as warnings signs of an early bubble.

From our vantage point, the timeline to humanoid maturation remains opaque, due to numerous engineering challenges, immature supply chains, and evolving regulatory frameworks. That said, conversations with industry participants suggest tangible progress is underway. Industry consensus now views meaningful commercial deployments at roughly five years away—a sharp compression from just two years prior, when many viewed general-purpose humanoids as a distant novelty, potentially decades from viability.

Recent policy announcements in both the United States and China have further elevated the strategic stakes for AI and robotics, increasingly framed through the lens of national security. The United States' approach under the second Trump administration has remained largely “light touch,” emphasizing open innovation, private-sector leadership, and targeted export controls to maintain its leadership in frontier models and computing infrastructure. In contrast, China is deploying the full weight of its top-down, state-directed model to systemically entrench embodied intelligence across all facets of its economy to achieve technological self-reliance and offset several demographic headwinds.


When viewed through the dual imperatives of aging populations and the race for “infinite labor” via self-replicating embodied AI, the near-term necessity and longer-term geopolitical implications come into sharp focus. China already dominates in mass manufacturing with a vertically integrated ecosystem of component suppliers and raw material processing. In response, we expect additional policy announcements from the United States in the coming 12 to 24 months, including a national robotics strategy or executive order, to address these critical gaps. The net result for investors should be a meaningful reduction in the risk of a classic “boom-bust” cycle as national security imperatives provide an increasingly credible government backstop to the private sector.

China's Existential Posture: Humanoids to Combat Aging Demographics

Robotics have been a cornerstone of China's global manufacturing ambitions since the early 2000s. According to the International Federation of Robotics (IFR), China became the world's largest user of industrial robots in 2013. The “Made in China 2025” plan designated robotics and AI to a top strategic priority, and the 2017 New Generation AI Development Plan framed them as a national security imperative. These policies drove rapid domestic adoption: by 2024, China's operational stock of industrial robots reached a record 2.027 million units, with annual installations of 295,000 units (54% of global demand), of which nearly 60% are manufactured domestically (source: IFR).

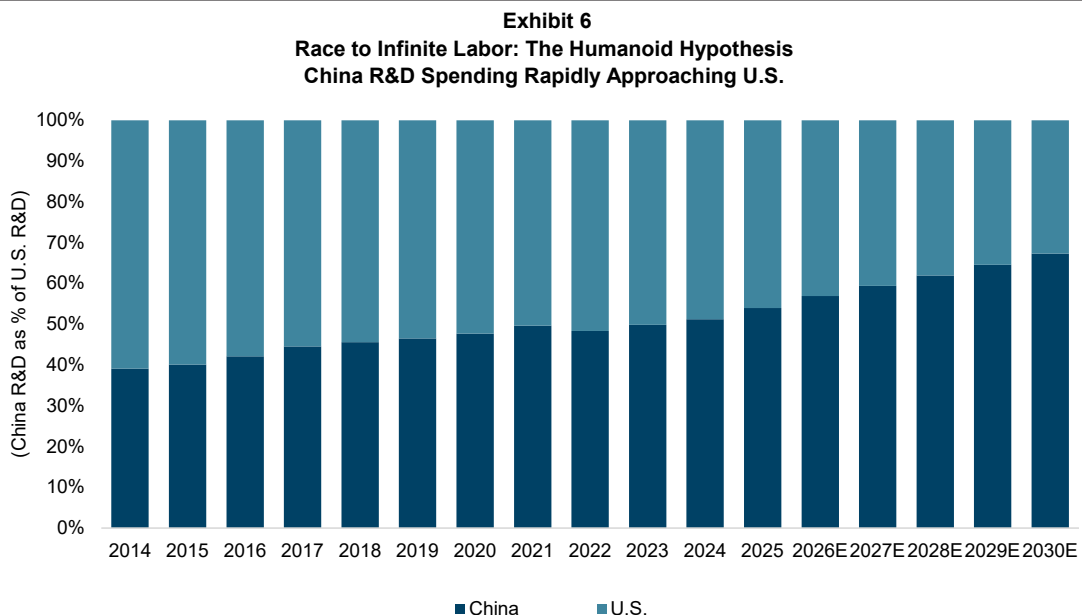
China's GDP growth has long outpaced that of the developed world but moderated since the end of strict COVID restrictions in early 2023, averaging around 4.3% to 5% in recent years; compared to roughly 5.5% in the United States over the same period. As the economy matures, Beijing is turning to innovation as a means of augmenting acute demographic challenges (aging population with low fertility rates) through sustained productivity. The 15th Five-Year Plan (2026-2030), adopted in March 2026, has seemingly adjusted to this new reality, targeting GDP growth of 4.5% to 5%,

while calling for R&D expenditure growth of at least 7%. At the current trajectory, we would expect China's R&D intensity to surpass 3% of GDP for the first time by 2030, with absolute R&D spending closing in on \$1 trillion by the early 2030s, as show in exhibit 5.



Sources: National Bureau of Statistics of China, William Blair Equity Research

Notably, on a purchasing power parity basis, China already exceeded \$1 trillion in R&D investment in 2024, in line with the United States' \$993 billion of combined federal and private-sector R&D spending. While the granular breakdowns for robotics and AI are not fully disclosed, the official targets point to substantial increases and potential convergence with the United States on an absolute basis in the coming decade, exhibit 6.



Sources: National Bureau of Statistics of China, National Center for Science and Engineering Statistics, William Blair Equity Research

Even more significant, and indicative of R&D priorities, China's policy announcements signal a full national mobilization in support of embodied AI and humanoids. In August 2025, the State Council released the "AI+" initiative, a comprehensive blueprint modeled on the highly successful "Internet+" campaign in 2015, which propelled China into the world's largest e-commerce market. The AI+ plan seeks to integrate AI across virtually every sector, targeting over 70% penetration of intelligent terminals and AI agents in key fields by 2027, rising to over 90% by 2030. A core objective is to create a flywheel of real-world data collection to accelerate the development of large language models, world models, and ultimately mass commercialization of humanoid robots and other embodied systems.

The 15th Five-Year Plan formally codifies the AI+ strategy and elevates embodied intelligence (including humanoids) as one of the core "future industries" alongside quantum computing and bio-manufacturing, among others. Top-priority designation delivers four key advantages:

1. Coordination across all levels of government;
2. Regulatory fast-track and preferential treatment;
3. Direct access to state-guided investment funds and subsidies; and
4. Performance accountability, with officials' evaluations tied to innovation KPIs.

Official projections call for AI-related industries to exceed CNY 10 trillion (\$1.45 trillion) by 2030, up from roughly CNY 1.2 trillion (\$170 billion) in core AI value-added output in 2025. This target is ambitious when compared with the United States, where the AI market had a comparable notional value in 2025 (measured by revenue) but is expected to grow at a more moderate private sector-driven pace, typically cited in the 20% to 30% CAGR range through 2030.


China has a proven track record of surpassing aggressive industrial goals in areas where scale, subsidies, and coordinated deployment provide an advantage (e.g., electric vehicles, industrial robots, telecommunications). However, it has faced difficulty in chokepoint technologies where high technical barriers, global supply chains, and limited substitutions come into play (e.g., semiconductors, high-end CNC, commercial aerospace). Humanoid robotics sit at the intersection of these crosscurrents: hardware price compression through mass deployment could accelerate progress, yet leading-edge computing power, high-performance actuators, and reliable real-world generation remain significant hurdles.

Underpinning this mobilization is a preemptive effort to strengthen dedicated capital for early-stage investments in emerging industries. In 2025, China launched the National Venture Capital Guidance Fund (targeting nearly CNY 1 trillion, or \$150 billion, over 20 years) and the National AI Industry Investment Fund (CNY 60 billion, or \$8.8 billion). Additional support comes from the mandated alignment of central state-owned enterprises (CITIC, SAIC, Sinopec, etc.), policy banks (e.g., Bank of China pledged CNY 1 trillion, or \$147 billion, through 2030), and provincial and municipal funds (such as Beijing, Shanghai, and Shenzhen committing tens of billions).

The VC Boom: U.S.’s Concentrated Bets Versus China’s Scattershot Approach

The venture capital market for robotics and embodied AI has experienced explosive growth since 2025, driven by heavy investments in humanoid manufacturers and component suppliers. Momentum accelerated further in early 2026, with capital inflows over the last 18-24 months surpassing the prior decade of cumulative robotics investment (2015-2024), as shown in exhibit 7.

Exhibit 7
Race to Infinite Labor: The Humanoid Hypothesis
Robotics and Embodied AI Venture Capital Funding



Note:


1) PitchBook search limited to 1) Humanoid Robotics, 2) Robotic Foundational Models, 3) Warehouse Robotics, 4) Robotic Process Automation.

2) The cited data has not been reviewed by PitchBook analysts and may be inconsistent with PitchBook methodology

Sources: PitchBook, a Morningstar company, William Blair Equity Research

The United States and China have emerged as the clear leaders in this space. Together they host roughly two-thirds of the world’s credible humanoid robotics companies and have accounted for over 80%-85% of global invested capital in the broader robotics and physical AI vertical since 2016, exhibit 8 below. This duopoly dynamic sits at the crux of the East-versus-West robotics arms race, as the exponential benefits of increased productivity look to impact all aspects of manufacturing and defense, and thereby, control of the geopolitical landscape in the decades ahead.

Exhibit 8
Race to Infinite Labor: The Humanoid Hypothesis
Cumulative Robotics and Embodied AI Venture Capital Funding



Note:

1) Using 2016 as base year

2) The cited data has not been reviewed by PitchBook analysts and may be inconsistent with PitchBook methodology

Sources: PitchBook, a Morningstar company, William Blair Equity Research

China's participation in the robotics and embodied AI funding cycle marks a departure from historical patterns that were long dominated by private and foreign investment. State-directed capital is now playing a significantly larger role contributing an estimated 50% of the deployed capital since 2025, compared to 25% to 40% in broader VC markets, historically. Rather than concentrating bets on a handful of companies, its strategy appears to reflect a deliberate scattershot approach, diversifying across dozens of competitors throughout different points of the supply chain. This approach creates intense domestic competition, accelerating real-world deployment and data generation, while maximizing shots on goal.

A distinctive feature has been the speed of financing; the last several years have seen leading companies secure multiple hundred-million-dollar rounds in quick succession. Flagship transactions include Galaxy General Robotics, or GalBot (multiple rounds totaling over \$800 million), Spirit AI Robotics (nearly \$600 million across rapid closings), TARS (\$455 million pre-A round), and substantial raises for AgiBot, LimX Dynamics, among others. This trend is now extending into the public markets, highlighted by Unitree Robotics' IPO filing for CNY 4.2 billion IPO (\$610 billion) in March 2026, with several others rumored to follow (Deep Robotics, GalBot, and AgiBot).

Whereas, the United States has pursued a classic bottom-up, private sector-driven model. Capital has been heavily concentrated on a smaller number of well-funded potential winners, such as Figure AI (\$2.3 billion raised), Apptrotronik (\$976 million), and Agility Robotics (\$683 million), and

often backed by Big Tech (e.g., Nvidia, Microsoft, Amazon, Google). While individual rounds have frequently been larger and accompanied by elevated valuations, the overall number of active humanoid players remains far more limited.

This concentrated approach allows deeper investment in advanced cognitive capabilities, safety systems, and high-value applications, but results in fewer parallel experiments and slower hardware iteration compared with China's broader ecosystem. The strategic contrast is stark, as China is betting on volume and speed versus the U.S. emphasis on quality and breakthrough innovation.


U.S. Playbook: Countering China Through Private Innovation and Strategic Guardrails

The United States has long maintained global leadership in semiconductor design, frontier AI models, and high-performance computing, but has lagged China in physical manufacturing at scale and supply-chain integration. The historical light-touch approach by the United States began to shift under the second Trump administration, which has framed advanced robotics as a national-security imperative and a direct counter to China's manufacturing dominance built over the past several decades.

The current administration's strategy can be viewed as a dual mandate: 1) a defensive pillar designed to restrict China's access to advanced semiconductors and AI chips, domestic markets, and domestic sources of capital, and 2) a proactive policy emphasizing deregulation, rapid infrastructure buildout, and financial incentives to encourage R&D and reshoring of critical supply chain (rare earth magnets, actuators, and high-end components). Key policy tools include:

- ***America's AI action plan*** (July 2025): A blueprint containing over 100 specific actions explicitly designed to win the AI and robotics race. It combines aggressive deregulation, accelerated infrastructure investment, international technology leadership, and targeted restrictions on adversarial access to U.S. innovation.
- ***Accelerating federal permitting of data center infrastructure: Executive Order 14318*** (July 2025): Streamlined federal permitting for large-scale data centers (projects over \$500 million or 100 MW), expanded access to federal lands, and introduced financial incentives such as loans, grants, and tax breaks.
- ***Genesis Mission: Executive Order 14363*** (November 2025): National effort, led by the Department of Energy to integrate AI with scientific discovery and advanced manufacturing, with an explicit focus on robotics and embodied intelligence.
- ***Ensuring a national policy framework for artificial intelligence: Executive Order 14365*** (December 2025): Establishes a uniform, minimally burdensome national AI regulatory framework and directs the Department of Justice to challenge overly restrictive state-level laws, effectively reducing compliance burdens on companies and accelerating the commercialization timeline for AI and humanoid technologies.

From a federal funding perspective, the U.S. has leaned heavily on its private sector, with only modest spending relative to China's state-guided funds, exhibit 9. Prior to 2020, dedicated federal robotics R&D was often measured in the tens of millions of dollars annually and embedded within broader AI or defense budgets. The National AI Initiative Act of 2020 (December 2020) marked a meaningful step forward, helping lift non-defense AI R&D and contributing to total federal AI-related investments rising to roughly \$3 billion annually in recent years. With the other notable standout being the Genesis Mission, which allocated \$293 million in fiscal 2026 for R&D in advanced manufacturing, autonomous systems, and embodied intelligence, and targeting \$1 billion through 2030.



Sources: National Center for Science and Engineering Statistics, International Monetary Fund, Whitehouse.gov, William Blair Equity Research

However, a significant portion of the increase in federal support is flowing through defense and national security channels, where robotics and embodied AI are increasingly viewed as critical dual-use technologies. The Department of Defense (War) and related agencies have steadily ramped up investment in autonomous systems, swarming technologies, and AI-enabled platforms, with the DoD launching an Artificial Intelligence Acceleration Strategy in January 2026. Recent budget priorities under this new initiative have shifted resources toward advanced manufacturing, unmanned ground vehicles, and humanoid-adjacent systems, with defense-related R&D budgets for emerging technologies now forming a growing share of the overall federal technology portfolio.

Exact figures for humanoid-specific programs remain classified or dispersed across multiple line items, although federal budget requests indicate a substantial increase in defense-related AI and autonomy spending since 2024, complementing the more-visible civilian efforts under the Genesis Mission. This combination of modest but targeted civilian funding and rising defense budgets reflects the administration's broader philosophy, using public resources to de-risk frontier technologies while letting the private sector drive scale and commercialization. Indicators of future policy announcements include:

- **Private sector meetings.** Commerce Secretary Howard Lutnick held closed-door meetings with humanoid and robotics CEOs in December 2025, suggesting administration support in accelerating the sector to specifically counter China. Notable humanoid companies in these discussions include Apptронik, Boston Dynamics, and Figure AI.
- **Bipartisan legislation.** The National Commission on Robotics Act (H.R. 7334, introduced February 2026) and the American Security Robotics Act (introduced March 2026 by Senators Tom Cotton and Chuck Schumer) would create an expert commission to assess competitiveness and explicitly prohibit federal procurement or funding of humanoid and unmanned ground systems from "covered foreign entities," e.g., China-based companies.
- **Robotics executive order.** The White House is rumored to be drafting an executive order on robotics expected in 2026, alongside early discussions of a national robotics strategy.

These measures look to extend the same national security logic that the United States has already applied to other physically integrated technologies. In 2019, President Trump signed Executive Order 13873, which established broad authority to block transactions involving information and communications technology from foreign adversaries. This authority was later used in January 2025 to impose a sweeping ban on Chinese and Russian hardware and software in connected vehicles (both vehicle connectivity systems and automated driving systems), as well as Chinese drone manufacturers.

While no sweeping bans on the importing of Chinese robotics or AI hardware have been formally announced, discussions, early legislative signals, and ongoing Section 232 investigations into robotics and industrial automation strongly indicate that additional policy announcements are likely in the next 12 to 18 months. These forthcoming measures are expected to provide further support of the domestic robotics and humanoid market through procurement preferences, targeted incentives, and supply-chain resiliency initiatives, while simultaneously expanding the ring-fence around Chinese systems in sensitive industries.

Major Bottlenecks to Scaling Humanoid Robotics

Supply-Chain Constraints and the Rush for Materials Security

Given the nascent stage of technology, global humanoid robot production was fairly limited in 2024, with shipments estimated in the low thousands of units (primarily pilots and prototypes). According to the International Data Corporation, shipments approximated 18,000 units in 2025, overwhelmingly led by Chinese manufacturers AgiBot, Unitree, and UBTECH Robotics. Leading U.S. and Chinese humanoid OEMs have outlined near-term production targets measured in the low tens of thousands of units in the coming years, with longer-term ambitions of scaling to hundreds of thousands annually, and in Tesla's case to over 1 million units.

The largest near-term challenge remains on the R&D front, as no humanoid OEM is widely viewed as having a fully viable model for mass commercialization yet. Once these technological hurdles are cleared, the bill of materials (BoM) is expected to meaningfully contract with scale. While disclosures are still limited, average BoM costs are expected to fall from roughly \$100,000 today to around \$20,000 to \$30,000, as annual production volumes exceed 10,000 units per platform.

However, U.S. manufacturers must also overcome significant supply chain challenges. Tesla has repeatedly highlighted actuators, motors, gears, and power electronics as the core hardware bottleneck to scaling. China dominates the manufacturing and processing of most of these components and the rare earth materials that account for the majority of a humanoid's BoM costs and performance. This leaves Western OEMs with a difficult choice: 1) continue to rely on sourcing from China with the risk of ongoing export controls, or 2) invest in vertical integration and Western suppliers. Key supply-chain bottlenecks include:

- **Actuators and joints:** These are the single largest cost and performance driver. Tesla has stated that actuators account for roughly 56% of Optimus total BoM cost. A typical Optimus Gen 2 requires approximately 28 actuators (14 linear and 14 rotary), with Gen 3 adding an additional 50 actuators for the hands alone. U.S. OEMs source approximately 50% to 70% of actuators and related components from Chinese suppliers (e.g., Ningbo Tuopu Group for linear and rotary assemblies, Zhejiang Sanhua for joint modules), although efforts are underway to onshore more proprietary designs.

- *Reducers (harmonic drives and planetary roller screws)*: These are critical subcomponents within actuators. Harmonic/strain wave drives are primarily used in rotary actuators for high-torque, zero backlash motion in shoulders, hips, waist, and elbows. Planetary roller screws are used in linear actuators for high-force, shock-absorption applications in weight-bearing joints such as knees, ankles, and calves.
- *NdFeB permanent magnets*: The strongest commercially available permanent magnets are essential for compact, high-torque motors inside actuators. Each humanoid requires approximately 2-3.5 kg of NdFeB magnets. These typically contain neodymium (Nd) along with small amounts of rare earths, dysprosium (Dy) and terbium (Tb), for higher temperature resistance to demagnetization.
- **Batteries and power systems**: Tesla Optimus Gen 3 uses a custom 2.3 kWh battery pack and represents less than 15% of BoM costs. However, there is a heavy reliance of global lithium-ion supply chains, where China dominates through companies such as Contemporary Amperex Technology Co. and BYD.
- **Sensors (vision, force/torque, tactile, IMUs)**: These account for approximately 10% to 20% of BoM costs. China benefits from large volume-based costs advantages in cameras, force sensors, and tactile arrays, while the United States leads in high-end AI integration but with challenges in affordable, high-volume scaling.

As previously highlighted in [Rare Earth Minerals: Potential for Critical Mineral Sovereign Fund and Other Backing](#) by William Blair analyst Neal Dingmann, China currently controls 70% of the mining, 90% of separation, and 93% of worldwide rare earth oxide and high-performance permanent magnet (NdFeB production). After imposing technology export restrictions in December 2023, China introduced licensing restrictions in April 2025 for key heavy rare earths (including dysprosium and terbium). This was followed by broader October 2025 measures for additional rare earth restrictions (12 in total) and requirements for government approval for foreign-made products containing even trace amounts (0.01%) of Chinese rare earth. While not full bans, these licensing regimes remain opaque and case-by-case, creating ongoing uncertainty for Western manufacturers.

In direct response, the United States has rapidly moved to alleviate the bottlenecks through a series of targeted reshoring measures and ongoing negotiations with China. To the second point, October 2025 controls have been suspended through November 10, 2026. And the status quo appears unchanged following the May 2026 Beijing summit, with hopes for an extension when President Xi Jinping visits the White House, scheduled for September 24, 2026; see Neal Dingmann's [summit note](#).

The United States has accelerated efforts to build domestic and allied rare earth and magnet capacity, by invoking the Defense Production Act (DPA) Title III to enable direct equity investments, loans, and expedited permitting for critical mineral projects. The most impactful step came in July 2025, when the DoD formed a major public-private partnership with MP Materials, securing a 15% equity stake, providing substantial financing, and securing long-term offtake and price floor commitments to establish end-to-end magnet manufacturing. Similar equity and loan deals, supported by both the DoD and Department of Commerce through CHIPS and Science Act extensions, have followed with American Resources Corporation, Vulcan Elements, and USA Rare Earth.

Additional measures include Project Vault (a strategic critical minerals reserve) and the January 2026 Section 232 action on processed critical minerals and derivative products (including rare earth permanent magnets). These are complemented by allied coordination, such as the April 2026 EU-U.S. Action Plan for Critical Minerals Supply Chain Resilience, which seeks coordinated trade policies and price floor mechanisms.

The United States has targeted operational domestic magnet production at scale by the late 2020s, with MP Materials' expanded facilities expected to begin meaningful output by approximately 2028. However, achieving broad independence from China for high-volume humanoid components remains a multiyear process.

Absent near-term disruption in materials sourcing, the next two years (2026-2027) represent a critical proof-of-scale period for U.S. humanoid OEMs. With successful policy execution, domestic mining and magnet production could support a ramp-up into the tens to hundreds of thousands of units by the early 2030s, although investment in precision manufacturing and a robust MRO ecosystem will still be required for long-term competitiveness.

Data Dilemma: Breaking the Real-World Training Paradox

Supply chain resilience forms the physical foundation, the next major barrier to scalable humanoid deployment is the acquisition of high-quality training data. Humanoids face a classic chicken-and-egg problem, as robots need high-quality real-world data to become capable and safe enough for widespread deployment, yet meaningful real-world data can only be generated at scale once large numbers of capable robots are already deployed. Three primary data sources used to train humanoid robots and embodied AI models include:

- ***Internet scale (human videos and egocentric footage)***: represents the most abundant source, with billions of hours of publicly available content online. It excels at capturing broad semantic understanding, task sequencing, object affordances, and high-level physical intuition, enabling pretraining of foundation models that transfer useful priors (probabilistic preconceptions) to robots. However, it suffers from fundamental limitations such as perception ambiguities, embodiment mismatch between human to robot kinematics, and lack of tactile feedback. These challenges require computer vision pipelines for converting 2D videos into structured 3D robot-specific action labels suitable for imitation learning.
- ***Simulation and synthetic***: provides effectively unlimited volume with perfect repeatability at very low cost, allowing safe exploration of edge cases. It dramatically accelerates reinforcement learning and world-model development by compressing thousands of years of experience into days of compute. Despite these strengths, the persistent sim-to-real gap remains a core bottleneck, as simulators struggle to model physical dynamics (friction, sensor noise, unstructured environments). Training solely in simulation frequently experiences sharp performance declines upon hardware deployment, limiting standalone value for final commercialization.
- ***Real-world robot***: provides the highest fidelity and ultimately most valuable data source, delivering synchronized multimodal streams (vision, force and torque, tactile, proprioception) that capture true physics that are impossible to replicate in video or simulation. It enables direct imitation learning, behavior cloning, and rapid policy improvement, while creating a powerful data flywheel once fleets are deployed.

Real-world data is the decisive accelerant for reliable generalization and mass commercialization. However, scarcity (due to limited operating humanoids) and high procurement costs create a major bottleneck. High-quality teleoperation data collection, annotation, and review can take 8 to 12 hours per usable hour of demonstration, with early commercial humanoid pilots primarily functioning as expensive data collectors, as noted by Tesla and other leaders. Also, it remains uncertain how much data will be required to achieve the diversity levels needed for commercial-level generalizations, or “zero-shot”—i.e., the ability to navigate and complete new tasks in unscripted environments.

At a high level, a single, complex manipulation task is thought to require several hundred to thousands of hours of high-quality demonstration. Tesla has stated that Optimus will need at least 10 times the data volume used for its Full Self-Driving (already power by over 10 billion fleet miles).


This scale implies the average humanoid platform will ultimately require several billion hours of embodied, multimodal data to achieve near zero-shot generalization levels. To put this into context, the hyperscalers have trained frontier models on tens of trillions of relatively cheaper digital tokens per training run, equivalent to just a couple hundred thousand hours of real-world data.

Leading humanoid OEMs are addressing data collection through aggressive, multipronged strategies. Main channels include 1) internal factory deployments, the highest-quality means of collection, 2) teleoperation farms using VR, exoskeletons, and motion capture suits (e.g., Tesla pays operators up to \$48 per hour, with each shift yielding 4 usable hours on average, and 3) open-source partnerships and collaboration. Notable examples include Figure AI’s testing at BMW’s Spartanburg, SC plant, AgiBot’s release of 1 million trajectories (3,000 hours) called AgiBot World 2026, and Google DeepMind’s Open X-Embodiment launch in October 2023 of 1 million real robot trajectories, in conjunction with 21 global research institutions.

Current methods rely on hybrid approaches that intelligently combine all three data types. The initial pilot phase is thought to require only a small amount of high-quality real-world robot data for final grounding (generated in prototype labs), when combined with internet-scale video for semantic pretraining and simulation for massive volume and edge-case exploration. Reinforcement learning is further closing the gaps, such as layering this hybrid foundation onto open-source datasets or using AI agents to control humanoids in simulated environments.

Cost trends are also encouraging, as the fully loaded cost of high-quality demonstration data has fallen from roughly \$340 per hour in early 2024 to approximately \$118 per hour as of early 2026, according to industry benchmarks (SVRC Research). This has been enabled by cheaper teleoperation hardware (now sub-\$2,000), better operator training, and standardized pipelines for the end-to-end capturing, processing, and storage of data. At the current trajectory, economically viable pricing (sub-\$10 per hour to compete with unskilled labor) could arrive in the early 2030s, unlocking mass commercialization.

Exhibit 10
Race to Infinite Labor: The Humanoid Hypothesis
Cost of Humanoid Training Data



Note:
 1) Fully loaded average cost per hour to produce
 Sources: SVRC Research: State of Robotics 2026, William Blair Equity Research

Regulatory and Safety Certification Roadblocks

Even with major breakthroughs in hardware, supply chains, and training data, humanoid robots cannot be deployed at meaningful scale in Western markets until regulatory bodies establish clear standards and OEMs successfully meet rigorous safety, liability, and cybersecurity requirements. Unlike software or consumer electronics, physical robots that operate alongside or in direct contact with humans face uniquely high barriers, as no company can legally sell or operate them in factories, warehouses, eldercare, or homes without certification. China is expected to follow its own, potentially faster regulatory pathway, but Western markets will remain gated by strict, independent safety approvals for the near term.

The regulatory landscape involves a layered ecosystem of organizations. In the United States, the most important bodies include OSHA (workplace safety under the General Duty Clause), Underwriters Laboratories (product safety standards such as UL 3300 for service robots), and American National Standards Institute/A3 (the core industrial robot safety standard being R15.06-2025). A critical near-term requirement is the Nationally Recognized Testing Lab (NRTL) field testing, an OSHA-administered evaluation at the actual deployment site.

Internationally, ISO leads with ISO 10218 (industrial robots) and the emerging ISO 25785-1 (safety requirements for mobile, dynamically stable robots including humanoids). Third-party certifier TÜV Rheinland is one of the most respected global bodies for functional safety testing and international certification. Whereas, in the United States, A3 (Association for Advancing Automation) serves as the leading standards developer and integrator certification authority. Additional oversight comes from the National Institute of Standards and Technology (technical research and guidelines) and, for consumer applications, the Consumer Product Safety Commission or U.S. Food and Drug Administration. Liability and data privacy rules (GDPR/CCPA equivalents) further complicate deployment.

Exhibit 11
Race to Infinite Labor: The Humanoid Hypothesis
United States Regulatory Timeline by End-Market

End-Markets	Expected Timeline	Primary Regulatory Bodies	Regulatory & Operational Status
Defense	2026	DOW, DARPA, NDAA	Currently funding, testing, and early deployments
Warehouses & Logistics	2027	ANSI/A3, ISO 10218/TS 15066, TÜV Rheinland	Active pilots; most advanced commercial pathway
Manufacturing	2027 – 2028	OSHA, ANSI/A3, ISO 10218, UL 3300	Strong progress with sandbox deployments
Construction	2028 – 2030	OSHA, ANSI/A3, ISO 25785-1	Early stage; high physical risk environment
Retail, Food, Hospitality	2028 – 2030	UL, ANSI/A3, ISO 25785-1, FDA/USDA, State Regulators	Early discussion; customer-facing adds complexity
Healthcare	2028 – 2030	UL, ANSI/A3, ISO 25785-1, FDA / State Regulators	Data-intensive; requires extensive safety data
Residential	2030's	CPSC, UL, State Regulators	No dedicated pathway yet; privacy and liability laws

Sources: Association for Advancing Automation, International Organization for Standardization, UL Solutions, OSHA, William Blair Equity Research

The primary challenges are fragmentation across jurisdictions, lengthy certification timelines (typically 18 to 36 months per model), and high costs for testing with iterative redesigns. Global standards remain unharmonized, creating complex multimarket compliance risks. Within the United States, Agility Robotics is farthest along, passing the NRTL field approval in a live commercial environment (Digit model). Although no company has received broad commercial certification for unrestricted warehouse or factory deployments. The other leading U.S. OEMs are in various stages of precertification or pilot validation, where humanoids are being tested in controlled industrial environments to generate the real-world safety data necessary to support full certification.

From Prototypes to Millions: Milestones and KPIs to Monitor

Just a few years ago, mass commercialization of general-purpose humanoid robots was widely viewed as decades away, often projected at 15 to 20 years. Progress was constrained by several fundamental challenges: immature AI and software lacking real-world adaptability, severely limited training data for physical tasks, unreliable and costly hardware (particularly hydraulic actuators and rudimentary end-effectors), and unproven integration into dynamic industrial or human environments.

Milestones

This timeline has compressed dramatically. As of mid-2026, consensus points to meaningful industrial scale deployment in warehouses and factories within roughly 3 to 5 years. This acceleration has been driven by the rapid convergence of several key breakthroughs between 2023 and early 2026. Chief among them are end-to-end vision-language-action (VLA) models (i.e., Nvidia's Project GR00T, Figure's Helix, and Tesla's vision-based systems), which have enabled robots to interpret natural language instructions, generalize from limited demonstrations, and shift from scripted behaviors to more adaptive autonomy.

Complementary hardware advances include the transition to fully electric high-performance actuators with integrated torque sensing and high degrees of freedom, dexterous multi-fingered hands with advanced tactile sensing, and improved power efficiency. These were supported by scalable real-world data collection through teleoperation and imitation learning, enhanced sim-to-real transfer via tools like Nvidia Isaac Sim, and vertical integration for cost reduction in core components, particularly through the leveraging of China's automotive supply chain.

Building on decades of foundational work in bipedal locomotion and dynamic balance, this 2023-to-2026 convergence of software and hardware has triggered an acceleration in humanoid startups and pilot deployments. The next 12 to 18 months are expected to prove pivotal, as the launch of next-generation models and platforms move the industry out of early pilot phases into broader commercial validation. If these efforts succeed, humanoid robots could begin operating in real-world industrial applications at meaningful scale within the next two to three years—although in limited capacity, as the pace of data collection and incremental hardware improvements dictate the speed at which the industry advances its S-curve. Key KPIs to monitor for the path to mass commercialization include:

1. **Safety.** As detailed in the prior section, safety and regulatory approval remain the foundational gating factors in the West. Initial certifications under ISO 25785-1 for warehouse and factory use are anticipated in late 2026 or early 2027, following expected publication of the standard. This will serve as a major de-risking milestone for broader industrial rollout. Sustained near-zero incident and fall rates under unsupervised operation will be a critical green light for broader deployment, into collaborative environments and eventual residential applications.

2. **Fault fates, mean time between failures (MTBFs).** Industrial customers require extremely high uptime, with fault rates low enough to support over 99.9% availability per shift and MTBF targets of at least 20,000 to 60,000 hours for dependable capital deployment. Current pilot systems still face notable challenges. Platforms relying heavily on early-generation Chinese actuators often achieve MTBF around 10,000 hours in lab or light-duty testing. However, real-world performance under dynamic loads is considerably weaker, implying a roughly 50% chance of at least one actuator failure within just a few hundred operating hours. Advancements are expected over the next several years, with several suppliers progressing toward more robust commercial-grade reliability, including Tesla, Hyundai, and Harmonic Drive Systems.
3. **Task completion rates.** This metric best reflects real-world usefulness and readiness for scale. In controlled industrial tasks (e.g., part handling, palletizing), the benchmark for commercialization is sustained greater-than-99% success rates with cycle times that match or exceed human performance. Real-world benchmarks remain significantly lower today in unstructured environments due to variable lighting, novel objects, human collaboration, and edge cases. Notable progress has already been demonstrated in industrial pilots.
 - Figure 02 robots at BMW’s Spartanburg, South Carolina, facility achieved over 99% placement accuracy per shift across more than 1,250 operational hours (November 2025 conclusion of an 11-month deployment), successfully loading over 90,000 sheet-metal parts and contributing to the production of more than 30,000 vehicles.
4. **Human-to-bot ratio.** This KPI directly measures labor multiplication potential and supervision economics. Early deployments often require one human overseeing just 1 to 3 robots. The commercialization threshold lies at ratios of 1:5, 1:10, or better, with minimal intervention needed for fleet management. Progress toward successful “lights-out” or low-supervision operations in factories and warehouses will confirm that humanoids are evolving from expensive assistants into genuine workforce multipliers.
5. **Durability of bot (lifecycle and total cost of ownership).** Long-term reliability determines whether robots deliver positive ROI over multiple years. Key metrics include overall service life (target 5 to 8 years), actuator and joint longevity, battery cycle life, and annual maintenance costs (ideally 10% to 5% or less of unit cost). Current dexterous hands and dynamic components remain wear points. Significant gains in materials, design, and maintainability are still required to unlock adoption beyond early adopter industrial sites.
6. **Scaling and bill-of-materials cost reduction.** Cost compression is the ultimate accelerator for mass production and market penetration. As discussed on page 14, BoM costs are expected to decline materially with incremental scale (tens of thousands of units) through standardized components, manufacturability design, and supply chain efficiency. Reaching hundreds of thousands of units per year is considered feasible, but scaling to over 1 million units annually will create intense pressure on the supply chain. Key bottlenecks include long lead times for capital equipment (CNC machines, giga-presses) and critical components.
 - For example, 1 million humanoids would require roughly 15 million to 20 million planetary roller screws, roughly the equivalent to today’s global annual production across all end-markets. While suppliers are beginning to scale, particularly in China, they are also cautious, carefully evaluating which platforms will emerge as winners before committing capital. The evolving technical landscape adds further complexity, as component requirements can shift with each design iteration (e.g., rotary versus linear actuators).

Framing the Opportunity

We view the labor market as an appropriate starting point for sizing the opportunity within humanoid robotics opportunity. Humanoids are designed to be biomimetic to humans and provide a drop-in replacement for any tasks humans can perform, which makes the relevant benchmark the portion of human labor that is susceptible to physical automation. In this context, the total addressable market should be understood as the economic value of work that could be augmented or replaced by humanoids over time.

We structure the analysis in three steps. We defined the relevant labor pool using occupational data, screened and ranked occupations based on humanoid functionality overlap, and then approximated employment wages to estimate the total addressable market.

For the United States, we rely on occupational employment and wage statistics from the U.S. Bureau of Labor Statistics. This dataset provides employment counts and average wages across occupations and offers a baseline for sizing labor markets. Not all occupations are equally exposed to humanoid automation, so we ranked them on five criteria: repetition, physicality, environment, dexterity, and human dependency.

The most exposed occupations are logistics and material movement, manufacturing and assembly, cleaning and maintenance, and food preparation. Moderately exposed occupations include construction labor, maintenance and repair, and certain agricultural roles. The least exposed occupations are professional services, management, education, and other knowledge-based roles that rely more heavily on cognitive work. The United States contains approximately \$1.3 trillion to \$1.7 trillion for annual wages in occupations with meaningful physical exposure (aggregated end-markets are displayed in exhibit 12). Although, the medium-term opportunity is measured in the \$3.8 trillion to \$5.4 trillion range, as technological advancements and regulations expand use-cases over time.

Exhibit 12
Race to Infinite Labor: The Humanoid Hypothesis
U.S. Market Opportunity by End-Market

End-Markets	Annual Total Labor Compensation (USD per worker)	Addressable U.S. Workforce	Tangible Addressable Market (USD)
Warehouses & Logistics	\$55K – \$70K	~6.5M – 7M	\$360B – \$490B
Manufacturing	\$75K – \$95K	~12.8M	\$960B – \$1.2T
Construction	\$60K – \$80K	~8.3M	\$500B – \$660B
Retail, Food, Hospitality	\$45K – \$60K	~15M – 18M	\$675B – \$1.0T
Healthcare	\$70K – \$100K	~18M – 20M (inc. healthcare support)	\$1.3T – \$2.0T

Sources: Bureau of Labor Statistics, William Blair Equity Research

We also look to the automotive market as another way of triangulating the opportunity ahead. We expect annual humanoid manufacturing to achieve, and eventually exceed that of passenger and commercial vehicles, as more derivative models for niche use-cases are developed in the ensuing decades. Using vehicle manufacturing as a proxy implies a nearly \$2.0 trillion opportunity,

assuming \$20,000 per unit versus global automotive production of 96.4 million units in 2025. We also highlight that China manufactured 34.5 million vehicles in 2025, more than the United States, Europe, and Japan combined. This is a staggering figure considering China's limited involvement entering the turn of the century and showcases its deep manufacturing moats.

Ways to Play Humanoid Robotics

There are few publicly traded humanoid OEMs today. Notable standouts including electric vehicle manufacturers, Tesla and Xpeng; Boston Dynamics' parent company, Hyundai Motor Group; and the sole pure play, UBTECH Robotics. The recent wave of venture capital investments in China is accelerating toward public markets, with Unitree Robotics having filed for an IPO in March 2026, with others advancing IPO preparations or guidance; namely AgiBot, Deep Robotics, and Leju Robotics. However, the leading Western OEMs continue to rely primarily on private capital, raising hundreds of millions in recent rounds often supported by strategic partners and corporates.

On the other hand, the "picks and shovels" approach offers investors early and diversified public-market exposure to the coming humanoid cycle. While current demand for hardware components remains nascent, the opportunity set among established suppliers is significantly more robust than among OEMs. Several key hardware architecture debates remain unresolved, making it challenging to identify long-term winners at this stage. Moreover, nearly all leading humanoid OEMs are pursuing some degree of vertical integration, viewing in-house design and manufacturing of core systems as essential to driving down BoM costs at scale.

A humanoid comprises dozens of component categories, from structural frames and wiring harnesses to batteries, sensors, and actuators. Many of these are expected to commoditize over time, or be sourced from established automotive, industrial, or consumer electronic suppliers. However, the highest value supplier opportunities are clustering around critical pain points where performance, reliability, and thermal constraints are severely constrained.

Humanoids are highly inefficient today, with thermal waste levels typically reported around 90%, meaning only 10% of electrical energy input converts into useful mechanical work. The rest is dissipated as heat, accelerating wear on components already subjected to repetitive, high-torque duty cycles. This is particularly acute in the actuators and motion control systems, where specialized manufacturing expertise and capital intensity act as a major barrier for full OEM insourcing.

This dynamic creates meaningful opportunities for public suppliers with deep expertise in high-temperature materials and alloys, advanced heat treatments, precision reducers, integrated cooling solutions, and high-tolerance components. While OEMs' aggressive BoM cost targets have raised investor concerns about potential price compression on suppliers at scale, we expect a multilayered supply chain to emerge. Where we anticipate similar characteristics of both the highly commoditized automotive supply chain and the higher-margin, proprietary aftermarket seen in commercial aerospace.

Exhibit 13
Race to Infinite Labor: The Humanoid Hypothesis
Humanoid Component Supply Chain

Component	Component Types	Level of Vertical Integration	Proprietary vs. Commoditized	Bottleneck Risk
Integrated Actuators (Comprises motors, gears, reducers, screws, and sensors)	Linear Actuators, Rotary Actuators, Direct Drive Actuators	High with U.S. OEMs; Mixed Elsewhere	Mostly Proprietary	High
Bearings, Gears, Reducers, Screws	Roller Bearings, Ball Bearings, Planetary Gears, Harmonic Reducers, Cycloidal Reducers, Ball Screws, Roller Screws	Low	Proprietary	Very high
Force and Torque Sensors	Single Axis Torque Sensors, 6-Axis Torque Sensors, Tactile Sensors	Mixed	Mostly Proprietary	Moderate to High
Motors	Frameless Brushless DC Motors, Permanent Magnet Synchronous Motors, Slotless Motors, Servo Motors	Mixed	Mostly Commoditized	Moderate
Hands/End Effecters	Tendon Driven Hands, Multi-Finger Hands, Grippers	High	Proprietary	Moderate
AI Compute	Fully Integrated System-on-a-Chip (SoC), GPUs	Low; Tesla is the exception	Proprietary	Moderate (NVIDIA supplier concentration risk)
Cameras/LiDAR	RGB Cameras, Depth Sensors, LiDAR	Low	Commoditized	Low
Batteries & Power Management	Battery Cells, Battery Management Systems	Low	Commoditized	Low
Motion Control Semiconductors	MCUs, Gate Drivers, Power MOSFETs, Encoders	Low	Commoditized	Low
Structure	Structural Frame, External Covering	Mixed	Proprietary Design; Commoditized Materials	Low

Sources: William Blair Equity Research

Exhibit 14
Race to Infinite Labor: The Humanoid Hypothesis
Select Humanoid Component OEMs and Relevant Commentary

Component OEM	Components Manufactured	Ticker	Country	Market Cap (in \$ Millions)	Humanoid Commentary
Schaeffler AG	Integrated Actuators and Subcomponents	SHA0-DE	Germany	\$10,455	"When we look at the bill of material of an average humanoid build for different purposes, we're talking about a 50% addressable market for Schaeffler ... Our best estimate at the moment is that this order book in total, order intake from the five customer contracts included, has a value somewhere in a midsize three-digit million range ... we are engaging today with 45 different customers. Engaging means active conversations, of which 30 prototype orders have resulted. And from these 30 prototype orders, five contracts have been secured." - 1Q26 Earnings Call (5/7/2026)
Regal Rexnord Corporation	Integrated Actuators and Subcomponents	RRX	USA	\$13,366	"There's anywhere from 30 to 40 axes on a humanoid and we can build that whole system ... and the value is we sit down with the OEM and we're designing jointly with them to come up with a solution ... they're not really so interested in focusing on the hardware, and that's what we bring to bear ... if you believed exactly their forecasts of volume, we would need to triple the size of the Regal in the next five years." - Investor Conference (2/19/2026)
Timken Company	Integrated Actuators and Subcomponents	TKR	USA	\$8,336	"Current offering allows value-added participation across an estimated 25% to 30% of bill of materials" - Investor Presentation (5/20/2026) "The humanoid market itself, we are working on applications today. We have a small amount of revenue. We'd expect a good CAGR off that small amount of revenue, but it takes a few years of a good CAGR to even get up to \$10 million. So, I think it's going to be a relatively small number the next couple of years." - 2Q25 Earnings Call (7/30/2026)
SKF AB Class B	Precision Bearings	SKF.B-SE	Sweden	\$11,649	Noted humanoid exposure, no direct commentary on customers or opportunity size
RBC Bearings Incorporated	Precision Bearings	RBC	USA	\$17,714	"Today we supply bearings for robotics that are in some pretty sophisticated applications where there is either high temperature or vacuum or a little bit of both producing computer chips. And the way this always starts is somebody that's designing a robot doesn't have any production volume and will go to one of our distributors and buy one of our bearings and use it in their prototype. And once it proves out and they start getting into production, they'll continue to use that distributor until production gets to a certain rate and then they'll trace back to the manufacturer or we'll find out about it from our distributor that this is an OEM that's using considerable amount of volume. And that's how all of these—every one of these markets is developed." - Fiscal 2Q26 Earnings Call (10/31/2025)
Harmonic Drive Systems Inc.	Gears & Reducers	6324-JP	Japan	\$4,171	"This year, we have finally taken the step to mass production to respond to the inquiries from the company, which is becoming from venture to manufacturer, in the North America ... The most prominent thing is that although we have not started mass production yet, we will definitely be approached when it comes to ultra-small and compact products. Also, in the upper limbs, arms, and shoulders, I would say we have done well ... The one North American company I mentioned earlier is definitely coming up in a plan. However, since this is still in the planning stage, we will need to make a more accurate assessment, but at least the scale of production in 2026 will be about three times that of 2025." - Fiscal 2Q26 Earnings Call (11/19/2025)
Nabtesco Corporation	Gears & Reducers	6268-JP	Japan	\$4,047	"Our Monocrank™ Series of Compact Precision Reduction Gears boasts the high rigidity, load capacity, and reliability cultivated through our RV™ Precision Reduction Gears as well as a compact, lightweight design. These gears demonstrate their unique characteristics when used for the shoulder and hip joints of humanoid robots, thereby enabling consistent, slight movements." - Investor Presentation (2/4/2026)

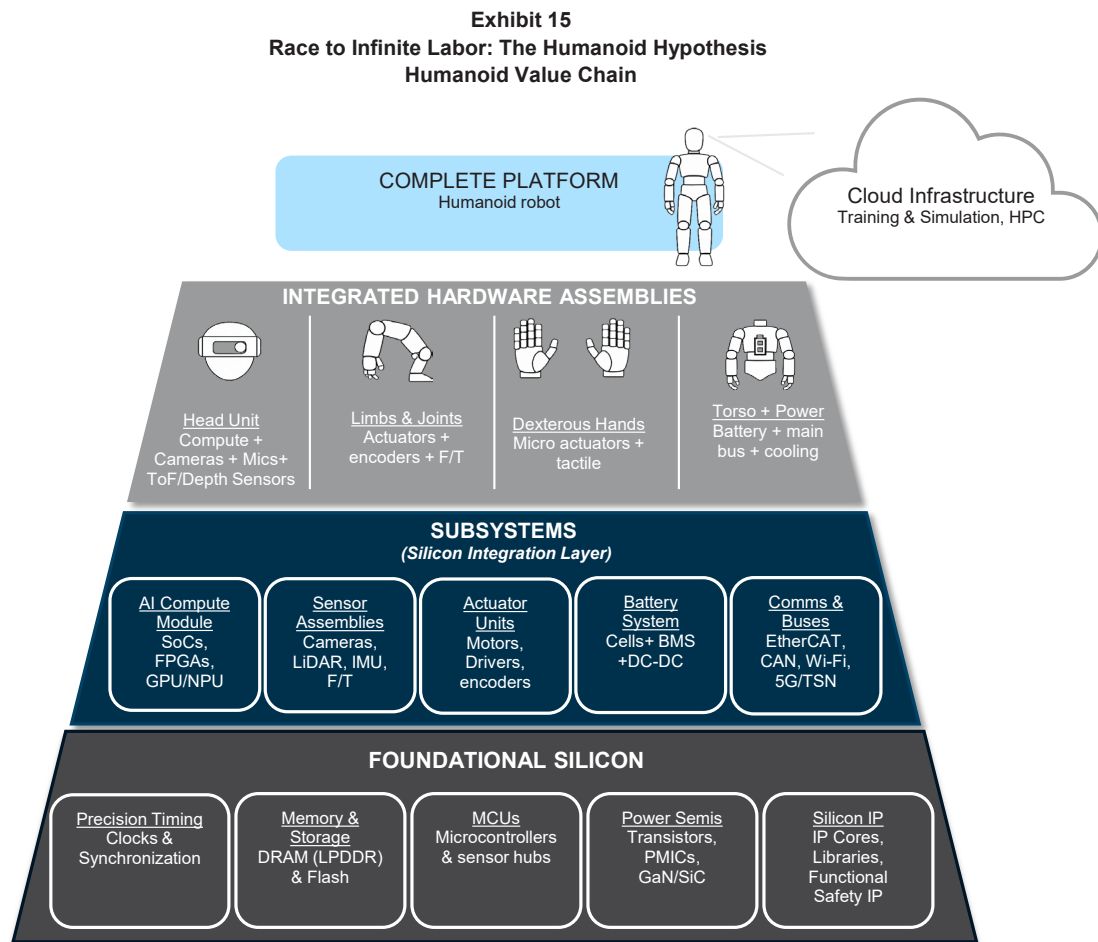
Exhibit 14 (continued)
Race to Infinite Labor: The Humanoid Hypothesis
Select Humanoid Component OEMs and Relevant Commentary

Component OEM	Components Manufactured	Ticker	Country	Market Cap (in \$ Millions)	Humanoid Commentary
THK Co., Ltd.	Ball and Roller Screws	6481-JP	Japan	\$5,110	Noted humanoid exposure, no direct commentary on customers or opportunity size
HIWIN Technologies Corp.	Ball and Roller Screws Motors	2049-TW	Japan	\$4,229	Noted humanoid exposure, no direct commentary on customers or opportunity size
NSK Ltd.	Ball and Roller Screws Precision Bearings	6471-JP	Japan	\$3,663	Noted humanoid exposure, no direct commentary on customers or opportunity size
Novanta Inc	Motors and Force/Torque Sensors	NOVT	USA	\$5,687	<p>"Novanta has unique capabilities that enable the perception or reaction of precision robotics in this physical world and to do so safely. We are excited about our recent design wins in the warehouse automation space, our momentum in surgical robotics and the ongoing development with multiple humanoid and warehouse automation players. We believe these physical AI applications are an important growth platform for Novanta, representing an incremental \$1 billion of addressable market by 2030." - 3Q25 Earnings Call (11/4/2025)</p> <p>"We're collaborating with over 10 leading humanoid players across North America and Europe. We're working with these players on multiple slots with significant content." 2Q25 Earnings Call (8/5/2025)</p>
Allegro MicroSystems, Inc.	Motors and Force/Torque Sensors	ALGM	USA	\$8,557	<p>"As humanoid robots are adopted in factories or in homes, a new rapid growth SAM will emerge for Allegro. Experts estimate the size of this robotics SAM to exceed \$10 billion between 2030 and 2035." - Fiscal 4Q25 Earnings Call (5/8/2025)</p> <p>"We secured two design wins with leading Chinese robotic companies for use in robotic joints during the quarter ... Initial shipments will begin in calendar 2026, with volumes expected to increase in 2027." - Fiscal 4Q26 Earnings Call (5/7/2026)</p> <p>\$150 of Addressable Content per Advanced Humanoid - Investor Presentation</p>
MinebeaMitsumi Inc.	Motors and Force/Torque Sensors Precision Bearings	6479-JP	Japan	\$10,076	"As for sensing devices, which I will touch on again later, we are seeing a rapid increase in inquiries for sensors equipped with strain gauges used in humanoid robots, for example sensors mounted on robot hands. We expect this area to gradually contribute to earnings going forward." - Fiscal 3Q2026 Earnings Call (2/5/2026)
Vishay Precision Group, Inc.	Motors and Force/Torque Sensors	VPG	USA	\$1,513	"We continue to make good progress across our business development initiatives, including humanoid robots and semiconductor equipment. These efforts generated \$11.8 million in orders during the fourth quarter, bringing total orders from these initiatives to \$37.8 million for the full year of 2025, which exceeded our goal of \$30 million for the year." - 4Q25 Earnings Call (2/11/2026)

Sources: Company Materials and Transcripts, FactSet, William Blair Equity Research

Layers of Humanoid Technology Stack

One way of framing humanoid robotics is simply transplanting AI into the physical world. This requires numerous different layers of technology working together to implement actions that resemble humans. As shown in the exhibit below, while the top layer represents a finished product that can perform real-world tasks, it is built on top of layers of intelligence, control, interface, and silicon infrastructure.



Sources: William Blair Equity Research

Foundational Infrastructure

The infrastructure layer is the foundation that enables every other layer to function reliably. Its role is not to compute or decide, but to ensure all compute, control, and sensing elements operate within tight electrical, thermal, and latency constraints. This includes converting battery power into usable voltage rails, distributing it efficiently, synchronizing components with precise clocks, and enabling communication across subsystems.

Key product categories include power semiconductors (PMICs, DC/DC converters, PoL regulators from Monolithic Power Systems, TI, Infineon), connectivity chips (Ethernet PHYs/switches and SerDes from Broadcom, Marvell, Credo, Astera Labs), and timing devices (MEMS oscillators and clock ICs from SiTime and Microchip). The layer is highly fragmented but mission critical, with bottlenecks increasingly shifting here as power density and data throughput scale.

Memory is a key constraint throughout the system, but for different reasons than within the data center (model size and complexity); robotics models add pressure by processing time-based, multimodal data. On the robot, memory is far more limited typically using LPDDR for main memory, SRAM for fast local access, and flash for storage. Power and space constraints on humanoids prevent data-center-style scaling, forcing a focus on efficiency driving model compression, reduced precision, and optimized memory use.

Physical Interface Layer

The physical interface layer is where the robot touches and perceives the real world converting digital commands into motion and physical signals back into data. Its function is bidirectional: actuation (motors, joints, end-effectors executing movement with force and precision) and sensing (vision, position, force, and environmental data). This is a physics-constrained layer where performance is dictated by torque density, response time, accuracy, and reliability rather than compute throughput, ultimately defining much of the robot's dexterity, smoothness, and perception quality.

Key products span actuation systems (servo motors, harmonic drives, gearboxes from Nidec, Nabtesco, Harmonic Drive Systems, and Yaskawa) and sensor stacks (image sensors from Sony, depth cameras from Intel RealSense, IMUs from Bosch, lidar from Velodyne/Ouster). On the semiconductor side, motor drivers, gate drivers, and sensor ICs (TI, Infineon, ADI) sit between control logic and physical components. FPGAs (AMD/Xilinx, Intel/Altera, Lattice) are often used for low-latency sensor processing and deterministic interfacing. CEVA's DSP and sensor fusion IP can be embedded in sensor hubs or edge processors that preprocess data before it moves up the stack.

The Control Layer

The control layer is the deterministic "nervous system" of the robot, turning high-level intent into precise, time-critical actions. It runs closed-loop control (PID loops, trajectory planning, safety checks) that continuously adjusts motor commands based on sensor feedback at millisecond or microsecond intervals. Unlike the intelligence layer, which is probabilistic and latency-tolerant, the control layer must guarantee hard real-time performance, stability, and safety.

Key products include MCUs, real-time processors, and motor control SoCs from STMicroelectronics, NXP, Renesas, TI, and Infineon, alongside real-time communication chips and protocols (Ethernet-CAT, CAN FD, industrial Ethernet PHYs) that synchronize distributed control across joints. FPGAs (AMD, Intel, Lattice) provide reconfigurable, ultra-low-latency hardware for real-time loops, sensor preprocessing, protocol handling, and bridging heterogeneous subsystems. IP providers are highly relevant: most MCUs are built on Arm Cortex-M/R cores, while CEVA contributes DSP and sensor fusion IP used in control and feedback loops.

Software Layer

The software/runtime layer is the orchestration and abstraction layer connecting AI decision-making to real-world execution. It provides a standardized environment where heterogeneous hardware (GPUs, MCUs, sensors, actuators) can interoperate, handling task scheduling, communication, data flow, and hardware abstraction. It ensures AI models do not manage low-level control directly and control systems do not need to interpret high-level intent. It also enables simulation and digital twin environments for training, testing, and validation before deployment, representing the "glue" that makes the stack programmable and scalable across robot designs.

Key products include robotics middleware and frameworks like ROS 2 (open-source standard) and commercial platforms such as Nvidia Isaac SDK and Omniverse, Microsoft Azure Robotics, and Unity for digital twins. The layer also includes RTOSes (VxWorks, QNX, FreeRTOS) and containerization/orchestration tools adapted for edge robotics. Unlike other layers, this one is software-centric but still tightly coupled with hardware capabilities and constraints.

Intelligence Layer

The intelligence layer is the “brain” of the robot, responsible for interpreting the world and deciding what to do. It runs AI models for perception (vision, sensor fusion), planning (navigation, manipulation), and reasoning (task execution, interaction), producing high-level intent passed to the control layer. Unlike control, this layer is probabilistic, compute-intensive, and latency-tolerant, operating in tens to hundreds of milliseconds. It is where most differentiation in humanoid robots will emerge as better models translate directly into more capable, adaptable, autonomous systems.

Key products include AI SoCs, GPUs, and vision processors such as Nvidia Jetson (e.g., Thor), Qualcomm Robotics RB-series, and Ambarella CVflow, alongside higher-end compute from AMD and Intel. The layer may also incorporate dedicated AI accelerators or FPGAs (AMD/Xilinx, Intel/Altera) for inference acceleration or custom model deployment. On the software side, it runs foundation and task-specific models from OpenAI, Google, or in-house teams (e.g., Tesla). Many edge AI SoCs are built on Arm CPU cores, while CEVA contributes neural processing and DSP IP for vision and sensor fusion workloads.

The intelligence layer also includes a tightly coupled memory subsystem (LPDDR or HBM from SK hynix, Samsung, Micron), which is critical to AI performance as it determines model capacity, bandwidth, and overall efficiency. While the intelligence layer uses the majority of humanoid memory, other components like MCUs, sensors, actuators also typically contain small amounts.

Mapping the Hardware and Silicon Content Inside a Humanoid

As shown in exhibit 16, on the following page, each humanoid has myriad different components manufactured by a diverse set of suppliers. Base silicon is built into larger hardware components to build motion, intelligence, and perception into each robot.

Exhibit 16 Race to Infinite Labor: The Humanoid Hypothesis Layers of a Humanoid

Layer 0: Semiconductor IP

CPU core IP	NPU / AI DSP IP	GPU / vision processing IP	Interconnect / NoC IP	Interface / PHY IP	Functional safety IP
SiFive Andes GlobalFoundries arm tenstorrent	Ceva cadence arm synopsys	Imagination arm	ARTERIS IP arm synopsys	cadence synopsys ALPHAWAVE SEMI Rambus	ARTERIS IP arm synopsys

Layer 1: Cloud Training & Simulation

AI Supercomputers/ training GPUs	AI Supercomputers/ training GPUs
NVIDIA intel AMD G	NVIDIA open robotics Google DeepMind unity

Layer 2: On-Robot AI Compute

Primary Compute Modules	FPGAs & Hardware Accelerators	Neuromorphic/Emerging
NVIDIA intel Qualcomm	AMD LATTICE SEMICONDUCTOR altera MICROCHIP	intel IBM brainchip*

Layer 3: Memory & Storage

DRAM	NAND Flash/NVMe SSDs	Industrial-Grade Memory & Storage	Memory IP & Interfaces
SAMSUNG micron SK hynix	SAMSUNG micron KIOXIA WD	swissbit* Apacer innodisk AEP	cadence synopsys Rambus

Layer 4: MCUs & Power Semis

Microcontrollers (Joint/Sensor)	Ultra-low-power MCUs & sensor hubs	Power semiconductors (IGBTs, GaN)	Power modules, DGDC, motor drivers
STMicroelectronics TEXAS INSTRUMENTS RENESAS NXP	infineon ambiq RENESAS	SILICON LABS TEXAS INSTRUMENTS onsemi	infineon ROHM SEMICONDUCTOR MPS VICOR ALLEGRO microsystems MAXLINEAR

Layer 5: Precision Timing

MEMS Oscillators & Resonators	Quartz & Legacy Timing	Clock Generators & PLL ICs
SiTime MICROCHIP	EPSON NIXX KYOCERA TXC	TEXAS INSTRUMENTS SKYWORKS RENESAS ANALOG DEVICES

Layer 6: Sensors & Perception

Cameras & depth sensors	LIDAR	IMUs & inertial	Force/torque & tactile	Event-based vision & encoders
intel StereoLabs* SONY OMNIVISION	OUSTER	STMicroelectronics BOSCH ANALOG DEVICES TDK InvenSense	ATI INDUSTRIAL AUTOMATION robot syntouch gelsight	PROPHESSEE RENISHAW HEIDENHAIN ami OSRAM

Layer 7: Actuators, Motors, and Mechanical

Motors & servo actuators	Reducers / harmonic drives	Dexterous hands / end effectors	Linear actuators & ball screws
MAXON Nidec Allient FAULHABER	Nabtesco	ROBOTIQ LEAP Hand (CMU) PSYONIC SCHUNK Shadow Robot	HIWIN NSK THK THOMSON

Layer 8: Power Systems

Battery Cells & Packs	BMS, power management & DGDC
SAMSUNG LG Energy Solution Panasonic	TEXAS INSTRUMENTS ANALOG DEVICES MPS VICOR NXP DELTA

Layer 9: Networking & Comms

On-Robot (Deterministic Buses)	External Wireless	Middleware/Robot OS
BECKHOFF IEEE BROADCOM	BOSCH MARVELL Qualcomm intel SEMTECH MEDIATEK	open robotics NVIDIA FOXGLOVE

Layer 10: Robot OEMs

TESLA Boston Dynamics* CLONE	SANCTUARY AI APPTRONIK	FIGURE AGILITY ROBOTICS ENGINEERED ARTS NEURA ROBOTICS	HUMANOID
------------------------------------	---------------------------	---	----------

Sources: William Blair Equity Research

Leading Humanoid OEMs

The humanoid landscape splits into two distinct philosophies. In the United States, most leading developers come from non-automotive roots with deep technology and AI backgrounds in the pursuit of high-performance, general-purpose platforms. Leading companies including 1X Technologies, Aptronik, Figure AI, and Tesla (see the appendix, on page 32, for a more detailed discussion) are targeting premium capabilities such as, advanced dexterity, human-like movement, and long-term versatility across factories, logistics, and eventually homes. These companies prioritize technical differentiation and are willing to accept higher initial costs to achieve superior performance and safety.

In contrast, Chinese players are moving faster and leaner, often with stronger automotive or consumer robotics backgrounds. XPeng and UBTECH bring direct automotive manufacturing expertise, while Unitree Robotics has emerged as the clear low-cost leader, deliberately optimizing for the “lowest viable” price point. Similarly, Agility Robotics in the U.S. follows a comparable cost-conscious strategy with its Digit platform, targeting immediate warehouse deployment rather than long-horizon general intelligence.

This creates a clear bifurcation of hyper-performance players (most U.S. leaders and XPeng) versus cost-optimized volume players (Unitree, Agility, and several Chinese newcomers). The former are betting on superior margins through premium capabilities, while the latter are racing to scale manufacturing and capture early market share through aggressive pricing and rapid iteration.

Exhibit 17
Race to Infinite Labor: The Humanoid Hypothesis
Major Humanoid OEMs

	Company	Headquarters	Strategic Investors ¹	Supply-Side Partnerships ¹	Demand-Side Partnerships ¹	AI Partnership ¹	Flagship Humanoid	Humanoid Description
North America	Agility Robotics	Salem, Oregon	Amazon, Schaeffler, NVIDIA, SoftBank	Schaeffler	Amazon, GXO Logistics, Schaeffler, Toyota, Mercado Libre	-	Digit	Digit features reverse jointed legs improving energy efficiency, load handling stability, and mobility while its simplified hand architecture improves the robot's reliability and cost profile.
	Aptronik	Austin, Texas	Google, Mercedes-Benz, Deere	Jabil	Mercedes-Benz, GXO Logistics, Jabil	Google	Apollo	Apollo features a modular architecture allowing customers to configure the robot in a number of form factors (bipedal, stationary, wheeled, movable arm, fixed arm) expanding the potential use-cases.
	Boston Dynamics	Waltham, Massachusetts	Hyundai, SoftBank	Hyundai	Hyundai	Google	Atlas	Atlas features 360-degree joint movement optimizing the robot for task efficiency across a number of industrial manipulation tasks, as well as advanced mobility functionality including running, jumping, and complex terrain navigation
	Figure AI	San Jose, California	Microsoft, OpenAI, NVIDIA, Jeff Bezos, Intel, Qualcomm	-	BMW, Brookfield	Microsoft, OpenAI	Figure	Figure 03 features the use of foam and textiles to soften the robot's exterior and advanced dexterity and motor controls to perform a variety of complex and variable tasks enabled by Figure AI's proprietary Helix vision-language-action AI model.
	Tesla	Austin, Texas	Publicly Listed	Tesla	Tesla	xAI	Optimus	Optimus is designed to be a multipurpose, widely deployable humanoid leveraging Tesla's vertically integrated supply chain.
	1X Technologies	Palo Alto, California	OpenAI, NVIDIA, EQT Group	-	EQT Group	OpenAI	NEO	Neo features a soft 3D-printed lattice polymer outer layer and a tendon drive system designed to mimic human musculature and improve safety in the event of a collision with humans and other obstacles.
China	AgiBot	Shanghai, China	Tencent, BYD, Baidu, LG Electronics	-	-	Baidu	A2	A2 features an expressive facial display screen designed for human interaction and advanced hand-level dexterity enabling small and unique object handling.
	Leju Robotics	Shenzhen, China	Tencent	-	-	Huawei	KUAVO	KUAVO 5 is a modular humanoid platform enabling customers to configure the robot in a number of iterations for specific tasks (bipedal, wheeled, hands, grippers, vision systems, sensor configurations).
	UBTech Robotics	Shenzhen, China	Publicly Listed	-	BYD, Dongfeng, Geely, FAW Group, Airbus	Huawei	Walker S2	Walker S2 is a designed to operate primarily in industrial manufacturing environments with advanced hand dexterity and load-bearing joints to improve lifting capability and durability.
	Unitree Robotics	Hangzhou, China	Tencent, Alibaba, Geely	-	-	DeepSeek	G1	The G1 is designed to be a low-cost and easy-to-deploy humanoid with the ability to perform basic tasks like moving objects and performing inspections.
	Xpeng Robotics	Guangzhou, China	Xpeng	Xpeng	Xpeng	Alibaba	Iron	IRON features a human-like endoskeleton, movement controlled by artificial muscles, and a soft exterior resulting in more human-like movement and appearance.
Europe & Middle East	Agile Robotics	Munich, Germany	SoftBank	-	-	Google	Agile ONE	Agile ONE features advanced arm and hand dexterity with a low center of mass optimized for stationary assembly and tool use.
	Hexagon Robotics	Stockholm, Sweden	Hexagon AB	Schaeffler, Hexagon AB, Maxon	Schaeffler, BMW	Microsoft	AEON	AEON features hybrid locomotion (wheels instead of feet) and advanced sensing enabling improved spatial awareness, 3D environment mapping, and enhanced inspection capabilities.
	Mentee Robotics	Tel Aviv, Israel	Mobileye	Aumovio	-	-	MenteeBot	MenteeBot is designed as a general-purpose humanoid designed to operate in dynamic human environments that runs on proprietary AI models designed to learn from and mimic human behavior.
	NEURA Robotics	Metzingen, Germany	Volvo	Schaeffler, Bosch, Omron	Schaeffler, Bosch	-	4NE1	4NE1 features advanced human-like hands and full body sensing capabilities enabling enhanced fine motor skills suited for complex tasks alongside human workers.
	RobCo	Munich, Germany	-	-	-	-	Alfie	Alfie's is designed as a stationary humanoid (comprising a torso, arms, hands/grippers, and a head) capable of completing repeatable industrial tasks including picking, kitting, and palletizing.

¹) Based on available public sources
Sources: Company materials, William Blair Equity Research

Appendix: Tesla's Optimus

Tesla's Optimus

Tesla's Optimus is seen as the Western leader in humanoids and currently the only public way for investors to get direct exposure to the industry. Tesla has numerous core competencies leveraged from its car business, including mass manufacturing, industrial robotics, AI autonomy software, training data centers, custom inference silicon, power electronics, and batteries. Optimus is the highest ROI embodiment of Tesla's vertically integrated real-world AI capabilities, and the company has increasingly presented itself as an AI robotics company rather than just an EV OEM, and we would argue the majority of its valuation is tied to its AI autonomy platform.

Tesla's Strategy


In Tesla's original debut of Optimus at AI Day in 2021, Elon Musk argued that Tesla's cars are already "robots on wheels," and that it therefore made sense to transfer the company's vision-based perception, inference, neural-network software, power electronics, battery systems, and actuator know-how into a humanoid form factor built for human environments. This strategic logic is Tesla's vertical-integration playbook, having repeated this in EVs by entering and innovating in areas where incumbents had established a complex tiered supply chain for batteries, electric motors, inverters, castings, charging, software, etc. Management now routinely links autonomy, robotaxi, Optimus, Dojo, and custom silicon as one integrated AI roadmap. In the company's AI and robotics materials, Tesla says the same end-to-end neural networks, onboard AI computer, and training infrastructure used in vehicles can be adapted to humanoids. In practical terms, that suggests Tesla wants Optimus to become three things at once: an internal manufacturing automation platform, a future external industrial labor product, and a broader physical-AI platform that may eventually move into household and service use-cases.

He also framed the initial use-case as eliminating "dangerous, repetitive, and boring" work. Today, Musk's proclamations have expanded considerably about the Optimus opportunity, calling it the most important product in humanity and that it will upend the economic function with such productivity that no one will have to work and it will usher in an era of abundance with high universal income.

Optimus Technical Review


Tesla originally provided deep technical detail about its progress but has since stopped updating the market because of IP and copycat concerns, especially from Chinese competitors. The debut at AI Day in 2021 described Optimus as a humanoid robot intended for human environments, moving at about five miles per hour, standing about 5'8" tall, and using eight cameras, Tesla's FSD computer, and the same perception tooling developed for vehicles. The stated goal was not just locomotion, but natural-language tasking in real-world settings.

Exhibit 18
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2021: Tesla Bot




Sources: Tesla

Exhibit 19
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2021: Technical Details



Sources: Tesla


Exhibit 20
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2021: AI for General Purpose Robotics



Sources: Tesla


At AI Day 2 in 2022, Tesla moved from concept art to hardware with a much deeper technical dive. Tesla's presentation described a production-intent design with a 2.3 kWh battery pack, a 52-volt electrical system, roughly 73 kg body mass, 28 structural actuators, and 11 degrees of freedom in each hand.

Exhibit 21
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2022: What's Changed



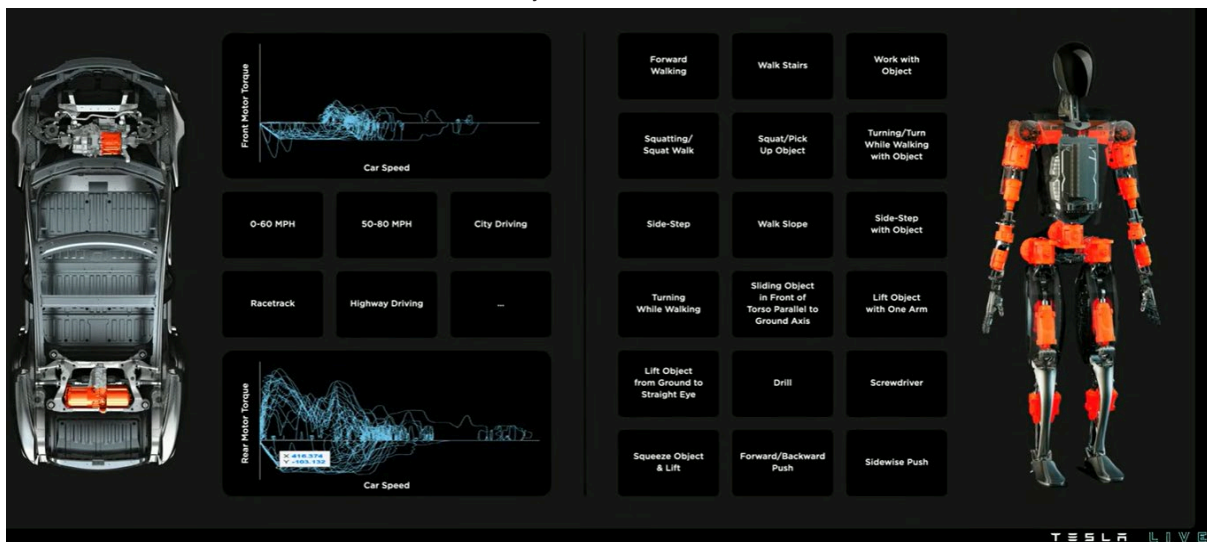
Sources: Tesla

Exhibit 22
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2022: Latest Generation



Sources: Tesla

Exhibit 23
Race to Infinite Labor: The Humanoid Hypothesis
Tesla AI Day 2022: Powertrain – Car vs. Robot



Sources: Tesla

In October 2024, Musk added that the “22 degrees of freedom” in the hands and forearms were the new state of art, and increased focus on the hand and wrist were the most important technical milestones.

The software architecture appears to be converging with Tesla’s vehicle stack. Tesla’s AI page says Optimus uses the same end-to-end neural networks, visual occlusion handling, AI inference computer, training infrastructure, and simulation capabilities used by the vehicle autonomy team.

The publicly shown milestones come at varying cadence, often via X posts. Tesla has now shown walking, balancing, squatting/yoga-like motions, object sorting, battery-cell handling, basic factory tasking, public untethered movement, and later household-task demonstrations such as cleaning

and simple manipulation. The second quarter 2024 shareholder deck said Optimus had begun performing tasks autonomously in one Tesla facility, and the second quarter 2024 earnings call said the robots were already performing tasks in the factory. Later public videos and event appearances suggested improving locomotion and upper-body control.

The long-awaited Optimus V3 was slated for April of 2026, but the event has been canceled and the production ramp-up delayed until year-end. Musk stated that there are more near-term revisions that need to be included in the hand actuators to improve object manipulation, and explained that each time Tesla gives a presentation, it finds competitors immediately reverse engineering. In other arenas like EV 48V architecture and super charging, Tesla has open-sourced its IP to push the industry forward, so this public protectionism is a new attitude and suggests heightened competition.

The unresolved technical questions remain the core investment debate. Tesla has not yet publicly disclosed robust data on mean time between failure, operating hours between interventions, task-success rates, service cost, joint wear, hand durability, grasp reliability, or performance in cluttered unstructured environments. Dexterity remains the key bottleneck for high-value non-repetitive work, and safety remains crucial because a humanoid robot operating around people has a much more demanding risk profile than a fenced industrial robot. The We, Robot event also underlined the need for caution: third-party reporting indicated that at least some showcased Optimus interactions relied on human assistance or teleoperation, which is not a disqualifier for development but does mean investors should not treat that event as proof of fully autonomous commercial readiness.

Exhibit 24
Race to Infinite Labor: The Humanoid Hypothesis
Optimus Development Timeline

Date	Event	Milestone	Commercial/Technical Significance
Aug. 2021	Tesla AI Day 2021	Concept Reveal	Human-scale robot using Tesla autonomy tools; prototype targeted “next year”
Oct. 2021	3Q 2021 Update	Early Interest	Tesla said AI Day drew “overwhelming interest” in Tesla Bot roles
Sep. 2022	Tesla AI Day 2022	Prototype Proof of Concept	Walking prototype onstage; production-intent architecture discussed
Jan. 2023	4Q 2022 Earnings Call	Strategy Articulation	Optimus linked to motors, batteries, power electronics, AI, and manufacturing
Jul. 2023	2Q 2023 Earnings Call	Custom Actuation	Tesla said no suppliers could provide adequate actuators; production-candidate internal actuators expected by Nov. 2023
Dec. 2023	Tesla Official Gen 2 Video	Optimus Gen 2 Announcement	Faster walking, better balance, improved hands and sensors
Jan. 2024	4Q 2023 Earnings Call	First Shipping Aspiration	Musk said Tesla had a “good chance” of shipping some units in 2025
Apr./Jul. 2024	2Q 2024 Earnings Call	Internal Factory Tasks	Tesla said Optimus began autonomous tasks in a facility and forecast several thousand units in Tesla factories by end of 2025
Oct. 2024	3Q 2024 Earnings Call	Hand / Forearm Dexterity Focus	Musk highlighted 22 DoF in hands and forearms
Oct. 2024	We, Robot	Public Showcase	Musk cited \$20k–\$30k eventual price range and called Optimus the biggest product ever
Jan. 2025	4Q 2024 Earnings Call	First Explicit Production Targets	“Roughly 10,000” internal plan for 2025; several thousand expected; outside customers maybe 2H26; cost <\$20k at >1M/yr
Apr. 2025	1Q 2025 Earnings Call	Factory Deployment	Musk said Tesla expected “thousands” of Optimus robots in Tesla factories in 2025
Jul. 2025	2Q 2025 Earnings Call	Public Visibility	Optimus walking in Palo Alto, serving popcorn at diner; prototype 3 expected within three months
Aug. 2025	Annual Meeting	Long-Term Ambition	Musk described production lines for Gen 3/4 and “tens of billions” long-term potential
Jan. 2026	4Q 2025 Earnings Call	Mass-Production design 1 / Fremont conversion	Tesla said Gen 3 is first mass-production design; first 1m/yr Optimus line by end of 2026 in Fremont; longer-term 10m/yr system discussed in Texas
Apr. 2026	1Q 2026 Earnings Call	Large-Scale Factory Preparation	Tesla said first large-scale line in Fremont designed for 1M robots/yr; Texas second-generation line aimed at 10M/yr long term

Sources: Tesla, William Blair Equity Research

The pattern of progress looks step-function rather than linear. Tesla’s public record shows long periods of conceptual framing followed by abrupt visible jumps, concept to walking prototype in 2022, custom-actuators in 2023, meaningful Gen 2 motion in 2023, internal factory task in 2024, and large-scale line planning in 2025–2026. That is encouraging from a technical-learning standpoint, but less reassuring for investors looking for predictable cadence of progress read-outs.

Exhibit 25
Race to Infinite Labor: The Humanoid Hypothesis
Commercialization Timeline and Production Commentary

Date	Speaker	Topic	Commentary
Aug. 2021	Musk	Initial Timing	Prototype “sometime next year”
Jul. 2023	Musk	Internal Deployment	Useful factory work “sometime next year”
Jan. 2024	Musk	Initial Shipment Ambition	“good chance of shipping some number” in 2026
Jul. 2024	Musk	Internal versus External Deployments	Several thousand in Tesla factories by end of 2025; external customers in 2027
Oct. 2024	Musk	Price Aspiration	We, Robot price range of \$20k–\$30k cited in press coverage
Jan. 2025	Musk	Production Target	“roughly 10,000” internal plan for 2026
Jan. 2025	Musk	Practical Revision	Several thousand likely by end of 2026
Jan. 2025	Musk	External Sales	Outside companies maybe in 2H27
Jan. 2025	Musk	Unit Cost	Cost <\$20k at >1M/yr
Jul. 2025	Musk	Long-Term Production Ramp	100,000/month over five years a “reasonable aspiration”
Aug. 2025	Musk	Production Ramp vision	Production line concepts ramping from 5k/month to 50k/month to 500k/month
Jan. 2026	Tesla / Musk	Factory Capacity	Fremont space converted into Optimus factory; long-term goal 1M/yr there
Apr. 2026	Tesla	Long-Term Capacity	Fremont first-generation line designed for 1m/yr; Texas second-generation line targeted for 10M/yr long term

Sources: Tesla, William Blair Equity Research

Musk is famous for his overly ambitious goals and time frames, and the commercialization timeline for Optimus is no different and has shifted over time. In 2021–2023, commentary centered on prototypes and “useful work.” In 2024, management shifted to internal factory deployment and specific dates for several thousand units. In 2025–2026, Tesla added harder numbers—10,000 robots in 2025 (was not achieved), 1-million-unit lines, and eventually 10-million-unit system design in Texas (annualized production figures). The helpful change is specificity. The unhelpful change is that the top-end volume ambition has expanded faster than the publicly validated installed base.

Terafab


Musk consistently looks at bottlenecks from first principles assuming everything goes perfectly to schedule, and if Optimus scales to 10 million units per year, Musk has determined that AI chips are the limiting factor. In Tesla fashion, the company is pursuing its most ambitious vertical integration yet, building a state-of-the-art 2 nm foundry to augment supply constraints from TSMC, Samsung, and Intel. Tesla and SpaceX have partnered in this endeavor and called it Terafab, with the plans

to build a massive AI chip manufacturing facility in Austin, Texas. Musk called it “the most epic chip-building exercise in history,” and recently announced it will be collaborating with Intel, likely adopting its 14A process. Long-term goals are to reach one terawatt of annual compute capacity, another signature aggressive objective Musk is famous for.

In the partnership, Tesla is first responsible for building and operating a research lab and pilot line with full end-to-end capabilities including design, lithography mask, fabrication, memory, logic, advanced packaging, and test under one roof. SpaceX will then be responsible for scaling up production when new chip designs are discovered and locked in, which we do not expect to be in full production before 2030.

Tesla’s in-house silicon design is called AI4 and is manufactured by Samsung. This inference chip is used in all its cars today for full self-driving (FSD) and in Optimus V3 prototypes. In April 2026, Musk stated that AI5 chip has reached tape-out and will be dual sourced with Samsung and TSMC, both providing a portion of the supply from their U.S. manufacturing facilities in Texas and Arizona, respectively. The question remains for investors, is chip production capacity in 2030 truly constrained? Tesla’s ambitious targets of 10 million Optimus units per year will each require an AI5 or AI6 chip.

Exhibit 26
Race to Infinite Labor: The Humanoid Hypothesis
Tesla Proprietary Silicon



Sources: Tesla

The prices of the common stock of other public companies mentioned in this report follow:

Advanced Micro Devices	\$467.51
Alphabet Inc. (Outperform)	\$379.38
American Resource Corporation (Outperform)	\$2.16
BYD Company	HK\$91.60
Contemporary Amperex Technology Co.	CNY 411.16
Harmonic Drive Systems Inc.	JPY 7,010
Hyundai Motor Company	KRW 655,000
Intel Corporation	\$119.84
Micron Technology, Inc. (Outperform)	\$751.00
MP Materials Corp (Outperform)	\$64.46
NVIDIA Corporation (Outperform)	\$215.33
Tesla, Inc. (Market Perform)	\$426.01
USA Rare Earth, Inc. (Outperform)	\$25.30

IMPORTANT DISCLOSURES

William Blair or an affiliate is a market maker in the security of RBC Bearings Incorporated, ON Semiconductor Corporation, Tesla, Inc., Micron Technology, Inc., Amazon.com, Inc., American Resources Corporation, MP Materials Corp., Novanta Inc., NVIDIA Corporation, STMicroelectronics N.V., Symbotic Inc. and USA Rare Earth, Inc.

William Blair or an affiliate expects to receive or intends to seek compensation for investment banking services from RBC Bearings Incorporated, ON Semiconductor Corporation, Tesla, Inc., Micron Technology, Inc., Amazon.com, Inc., American Resources Corporation, MP Materials Corp., Novanta Inc., NVIDIA Corporation, STMicroelectronics N.V., Symbotic Inc. and USA Rare Earth, Inc. or an affiliate within the next three months.

Officers and employees of William Blair or its affiliates (other than research analysts) may have a financial interest in the securities of RBC Bearings Incorporated, ON Semiconductor Corporation, Tesla, Inc., Micron Technology, Inc., Amazon.com, Inc., American Resources Corporation, MP Materials Corp., Novanta Inc., NVIDIA Corporation, STMicroelectronics N.V., Symbotic Inc. and USA Rare Earth, Inc.

This report is available in electronic form to registered users via R*Docs™ at <https://williamblairlibrary.bluematrix.com> or www.williamblair.com.

Please contact us at +1 312 236 1600 or consult <https://www.williamblair.com/equity-research/coverage> for all disclosures.

Ross Sparenblek, Sebastien Naji and Jed Dorsheimer attests that 1) all of the views expressed in this research report accurately reflect his/her personal views about any and all of the securities and companies covered by this report, and 2) no part of his/her compensation was, is, or will be related, directly or indirectly, to the specific recommendations or views expressed by him/her in this report. We seek to update our research as appropriate. Other than certain periodical industry reports, the majority of reports are published at irregular intervals as deemed appropriate by the research analyst.

DOW JONES: 50579.70
 S&P 500: 7473.47
 NASDAQ: 26344.00

Additional information is available upon request.

Current Rating Distribution (as of May 26, 2026):

Coverage Universe	Percent	Inv. Banking Relationships *	Percent
Outperform (Buy)	74	Outperform (Buy)	13
Market Perform (Hold)	25	Market Perform (Hold)	2
Underperform (Sell)	1	Underperform (Sell)	0

*Percentage of companies in each rating category that are investment banking clients, defined as companies for which William Blair has received compensation for investment banking services within the past 12 months.

The compensation of the research analyst is based on a variety of factors, including performance of his or her stock recommendations; contributions to all of the firm’s departments, including asset management, corporate finance, institutional sales, and retail brokerage; firm profitability; and competitive factors.

OTHER IMPORTANT DISCLOSURES

Stock ratings and valuation methodologies: William Blair & Company, L.L.C. uses a three-point system to rate stocks. Individual ratings reflect the expected performance of the stock relative to the broader market (generally the S&P 500, unless otherwise indicated) over the next 12 months. The assessment of expected performance is a function of near-, intermediate-, and long-term company fundamentals, industry outlook, confidence in earnings estimates, valuation (and our valuation methodology), and other factors. Outperform (O) - stock expected to outperform the broader market over the next 12 months; Market Perform (M) - stock expected to perform approximately in line with the broader market over the next 12 months; Underperform (U) - stock expected to underperform the broader market over the next 12 months; not rated (NR) - the stock is not currently rated. The valuation methodologies include (but are not limited to) price-to-earnings multiple (P/E), relative P/E (compared with the relevant market), P/E-to-growth-rate (PEG) ratio, market capitalization/revenue multiple, enterprise value/EBITDA ratio, discounted cash flow, and others. Stock ratings and valuation methodologies should not be used or relied upon as investment advice. Past performance is not necessarily a guide to future performance.

The ratings and valuation methodologies reflect the opinion of the individual analyst and are subject to change at any time.

Our salespeople, traders, and other professionals may provide oral or written market commentary, short-term trade ideas, or trading strategies to our clients, prospective clients, and our trading desks that are contrary to opinions expressed in this research report. Certain outstanding research reports may contain discussions or investment opinions relating to securities, financial instruments and/or issuers that are no longer current. Investing in securities involves risks. This report does not contain all the material information necessary for an investment decision. Always refer to the most recent report on a company or issuer. Our asset management and trading desks may make investment decisions that are inconsistent with recommendations or views expressed in this report. We will from time to time have long or short positions in, act as principal in, and buy or sell the securities referred to in this report. Our research is disseminated primarily electronically, and in some instances in printed form. Research is simultaneously available to all clients. This research report is for our clients only. No part of this material may be copied or duplicated in any form by any means or redistributed without the prior written consent of William Blair & Company, L.L.C.

This is not in any sense an offer or solicitation for the purchase or sale of a security or financial instrument.

The factual statements herein have been taken from sources we believe to be reliable, but such statements are made without any representation as to accuracy or completeness or otherwise, except with respect to any disclosures relative to William Blair or its research analysts. Opinions expressed are our own unless otherwise stated and are subject to change without notice. Prices shown are approximate.

This report or any portion hereof may not be copied, reprinted, sold, or redistributed or disclosed by the recipient to any third party, by content scraping or extraction, automated processing, or any other form or means, without the prior written consent of William Blair. Any unauthorized use is prohibited.

If the recipient received this research report pursuant to terms of service for, or a contract with William Blair for, the provision of research services for a separate fee, and in connection with the delivery of such research services we may be deemed to be acting as an investment adviser, then such investment adviser status relates, if at all, only to the recipient with whom we have contracted directly and does not extend beyond the delivery of this report (unless otherwise agreed specifically in writing). If such recipient uses these research services in connection with the sale or purchase of a security referred to herein, William Blair may act as principal for our own account or as riskless principal or agent for another party. William Blair is and continues to act solely as a broker-dealer in connection with the execution of any transactions, including transactions in any securities referred to herein.

For important disclosures, please visit our website at williamblair.com.

This material is distributed in the United Kingdom and the European Economic Area (EEA) by William Blair International, Ltd., authorised and regulated by the Financial Conduct Authority (FCA). William Blair International, Limited is a limited liability company registered in England and Wales with company number 03619027. This material is only directed and issued to persons regarded as Professional investors or equivalent in their home jurisdiction, or persons falling within articles 19 (5), 38, 47, and 49 of the Financial Services and Markets Act of 2000 (Financial Promotion) Order 2005 (all such persons being referred to as "relevant persons"). This document must not be acted on or relied on by persons who are not "relevant persons."

This report is being furnished in Brazil on a confidential basis and is addressed to the addressee personally, and for its sole benefit. This does not constitute an offer or solicitation for the purchase or sale of a security by any means that would constitute a public offering in Brazil under the regulations of the Brazilian Securities and Exchange Commission (*Comissão de Valores Mobiliários*) or an unauthorized distribution under Brazilian laws and regulations. The securities are authorized for trading on non-Brazilian securities markets, and this report and all the information herein is intended solely for professional investors (as defined by the applicable Brazilian regulation) who may only acquire these securities through a non-Brazilian account, with settlement outside Brazil in a non-Brazilian currency.

"William Blair" and "R*Docs" are registered trademarks of William Blair & Company, L.L.C. Copyright 2026, William Blair & Company, L.L.C. All rights reserved.

William Blair & Company, L.L.C. licenses and applies the SASB Materiality Map® and SICSTM in our work.

Equity Research Directory

John Kreger, Partner Director of Research +1 312 364 8612
Kyle Harris, CFA, Partner Operations Manager +1 312 364 8230

CONSUMER

Sharon Zackfia, CFA, Partner +1 312 364 5386
Group Head–Consumer
Lifestyle and Leisure Brands, Restaurants, Automotive/E-commerce

Jon Andersen, CFA, Partner +1 312 364 8697
Consumer Products

Phillip Blee, CPA +1 312 801 7874
Home and Outdoor, Automotive Parts and Services, Discount and Convenience

Dylan Carden +1 312 801 7857
E-commerce, Specialty Retail

ECONOMICS

Richard de Chazal, CFA +44 20 7868 4489

ENERGY AND POWER TECHNOLOGIES

Jed Dorsheimer +1 617 235 7555
Group Head–Energy and Power Technologies
Generation, Efficiency, Storage

Neal Dingmann +1 312 801 7835
Oil and Gas, Rare Earth

Tim Mulrooney, Partner +1 312 364 8123
Energy and Environmental Services

FINANCIAL SERVICES AND TECHNOLOGY

Adam Klauber, CFA, Partner +1 312 364 8232
Group Head–Financial Services and Technology
Financial Analytic Service Providers, Insurance Brokers, Property & Casualty Insurance

Andrew W. Jeffrey, CFA +1 415 796 6896
Fintech

Cristopher Kennedy, CFA +1 312 364 8596
Fintech, Specialty Finance

Jeff Schmitt +1 312 364 8106
Wealthtech, Wealth Management, Capital Markets Technology

GLOBAL SERVICES

Tim Mulrooney, Partner +1 312 364 8123
Group Head–Global Services
Commercial and Residential Services

Andrew Nicholas, CPA +1 312 364 8689
Consulting, HR Technology, Information Services

Trevor Romeo, CFA +1 312 801 7854
Staffing, Waste and Recycling

HEALTHCARE

Biotechnology

Matt Phipps, Ph.D., Partner +1 312 364 8602
Group Head–Biotechnology

Sami Corwin, Ph.D. +1 312 801 7783

Lachlan Hanbury-Brown +1 312 364 8125

Andy T. Hsieh, Ph.D., Partner +1 312 364 5051

Myles R. Minter, Ph.D., Partner +1 617 235 7534

Scott Hansen, Partner Associate Director of Research +1 212 245 6526

Healthcare Technology and Services

Ryan S. Daniels, CFA, Partner +1 312 364 8418
Group Head–Healthcare Technology and Services
Healthcare Technology, Healthcare Services

Steven Lichtman +1 646 561 9373
Medical Technology

Brandon Vazquez, CFA +1 212 237 2776
Dental, Animal Health, Medical Technology

Life Sciences

Matt Larew, Partner +1 312 801 7795
Life Science Tools, Bioprocessing, Healthcare Delivery

Andrew F. Brackmann, CFA +1 312 364 8776
Diagnostics

Max Smock, CFA +1 312 364 8336
Pharmaceutical Outsourcing and Services

INDUSTRIALS

Brian Drab, CFA, Partner +1 312 364 8280
Co-Group Head–Industrials
Advanced Manufacturing, Industrial Technology

Ryan Merkel, CFA, Partner +1 312 364 8603
Co-Group Head–Industrials
Building Products, Specialty Distribution

Louie DiPalma, CFA, Partner +1 312 364 5437
Aerospace and Defense, Smart Infrastructure

Ross Sparenblek +1 312 364 8361
Diversified Industrials, Robotics, and Automation

TECHNOLOGY, MEDIA, AND COMMUNICATIONS

Jason Ader, CFA, Partner +1 617 235 7519
Co-Group Head–Technology, Media, and Communications
Infrastructure Software

Arjun Bhatia, Partner +1 312 364 5696
Co-Group Head–Technology, Media, and Communications
Software

Dylan Becker, CFA +1 312 364 8938
Software

Louie DiPalma, CFA, Partner +1 312 364 5437
Government Technology

Jonathan Ho, Partner +1 312 364 8276
Cybersecurity, Security Technology

Patrick McIlwee, CFA +1 312 364 8655
Software

Sebastien Naji +1 212 245 6508
Infrastructure Software, Semiconductor and Infrastructure Systems

Maggie Nolan, CPA, Partner +1 312 364 5090
IT Services

Ralph Schackart III, CFA, Partner +1 312 364 8753
Internet and Digital Media

Stephen Sheldon, CFA, CPA, Partner +1 312 364 5167
Vertical Technology – Real Estate, Education, Restaurant/Hospitality

EDITORIAL AND SUPERVISORY ANALYSTS

Steve Goldsmith, Head Editor and SA +1 312 364 8540

Katie Anderson, Editor and SA +44 20 7868 4451

Audrey Majors, Editor and SA +1 312 364 8992

Beth Pekol Porto, Editor and SA +1 312 364 8924

Lisa Zurcher, Editor and SA +44 20 7868 4549