

THE VERIFIER - REPORTS

REPORT NO. 03



The Humanoid Economy

"General-purpose robots have crossed from demonstration to industry."

CLAIM UNDER TEST

JULY 2026

FREE TO READ - NEVER SPONSORED

SHA-256 3a7ca3a2fd1c232a13da07ebc7763c96
26bdc0e609dc9ae833b530fa227612d9

FINGERPRINT OF THE CANONICAL TEXT - RECIPE IN APPENDIX B

How to read this report	3
PART ONE - The Question	5
01 The claim under test	5
02 The Five Gates	6
03 The year the demos ended	7
PART TWO - The Gates	9
04 Gates one and two: autonomy, and the split	9
05 Gate three: the order book	9
06 Gate four: the utilisation void	10
07 Gate five: capacity, components and the wager	11
08 The rulebook arrives	12
PART THREE - The Economics	15
09 What a robot costs, and what a robot-hour costs	15
10 The verifier's toolkit	15
PART FOUR - The Scoreboard	18
11 Three years on the record	18
12 Twelve months of checkable markers	18
13 What would change our verdict	19
14 The verdicts	19
Appendix A - The claims ledger	20
Appendix B - Method and fingerprint	23
Appendix C - Sources	23

How to read this report

This is a Verifier report and it follows the house rules: one question, stated on the cover as a claim; every material claim logged in the ledger at Appendix A with a status and a source; a SHA-256 fingerprint of the canonical text printed on this page and recomputable by any reader using the recipe in Appendix B; a section naming what would change our verdicts; and a corrections log on the public record. It is the third in the series and the robotics desk's flagship: where Report No. 01 asked what can be verified about AI and No. 02 asked what proof costs, this report asks the older, harder question the desk was founded on - when a machine is claimed to work, what would it take to know?

The independence note is standard: no sponsor touches research, ever. The conflict note is short this time: the commercial roles disclosed in Reports No. 01 and 02 - in proving infrastructure and agent accountability - do not intersect the robotics category, so this report carries no coverage exclusions. It does carry a naming policy worth stating: this desk names names. A deployment claim is checkable only when the customer, the site and the workload are on the record, so companies appear here with their claims attached and their evidence graded, which is a compliment to the ones who publish numbers and a finding about the ones who do not.

One reading instruction, adapted from the series' standing warning about multipliers. In robotics the unit of self-promotion is not the speedup but the operating statistic - the tote count, the vehicle count, the hours logged - and nearly every such number in circulation is reported by the party that benefits from it, usually the vendor, occasionally the customer, almost never a third party. Where this report quotes one, the ledger records who counted. The desk's demo manual compresses the discipline to a sentence this report will spend twenty pages earning: a demonstration is a claim, an invoice is evidence, and a maintenance log is proof.

T H E F I N G E R P R I N T

SHA-256 3a7ca3a2fd1c232a13da07ebc7763c9626bdc0e609dc9ae833b530fa227612d9

Recompute it: take report-03-canonical.txt, decode entities, remove all whitespace, lower-case, hash. Full recipe in Appendix B.

PART ONE

The Question



The claim under test

The claim arrives from every direction at once. From the keynote stages, where humanoid robots now box, flip and run half-marathons; from the order books, where a Chinese maker's audited filing shows thousands of units shipped and tens of thousands of committed capacity; from the factory floors of three continents, where the world's most famous automakers publish pilot results with part-counts attached; and from the venture market, which has priced the leading Western maker in the tens of billions of dollars on the strength of it. The claim is that the general-purpose robot - the machine shaped like us, built to work where we work - has crossed from demonstration to industry.

This report tests that sentence against the record as of July 2026, and the finding is that the crossing is real, narrow, and radically uneven across the five things "industry" actually requires. There is now verified, third-party-visible productive work: a warehouse robot past one hundred thousand totes under a paying contract, a factory robot through eight-hour shifts for over a year, an automotive pilot with placement-accuracy statistics attached. There is a genuine order book, dominated in volume by China and in evidence-per-unit by the West. And there is an equally genuine void exactly where an industry's existence is normally proven: almost no maker publishes utilisation, reliability, intervention rates or the month-eighteen numbers that separate a machine that works from a machine that worked on filming day. The supply side of the humanoid economy exists beyond dispute - factories, filings, standards, service contracts. The demand side, measured in robots verifiably earning their keep, still numbers in the hundreds, worldwide, doing narrow tasks superbly counted.

THE SCOREBOARD - SIX SEGMENTS, ONE STANDARD

SEGMENT	WHAT THE RECORD SHOWS	STATUS
Demonstration capability	Athletic, dexterous, continuous - and a claims category, not an evidence category	VERIFIED as theatre
Autonomy	The industry's open secret is now its open dispute: teleoperation as business model versus autonomy as doctrine	CONTESTED
The order book	Audited thousands shipped in China; verified dozens working in the West; buyer mix still dominated by laboratories	TRACKING
Unit economics	Utilisation, reliability and attendant ratios almost nowhere disclosed - the void where the industry would be proven	EARLY
Capacity and supply chain	Committed capacity runs ten to twenty times shipments; components concentrate in one country	TRACKING - a disclosed wager
Rules and standards	The industrial rulebook just modernised; the walking-robot standard is still a draft	EARLY

Three numbers frame everything that follows. The first is one hundred thousand - the totes moved by a single robot model under a single paying contract, the most thoroughly verified productive work any humanoid has done, and a useful scale for every larger claim in this report. The second is seventy - the percentage of the volume leader's 2025 humanoid shipments that went to research and education buyers, by its own audited filing: the single most honest maturity gauge the industry has published, and it says laboratory. The third is zero - the number of published international safety standards written for a robot that can fall over, as of this report's date; the industrial rulebook was modernised last year for arms and applications, and the standard for dynamically stable machines is still in the drafting stage while those machines ship to factories and living rooms.


The supply side of the humanoid economy is beyond dispute. The demand side, measured in robots verifiably earning their keep, numbers in the hundreds - doing narrow work, superbly counted.

CHAPTER 02

The Five Gates

Industries are not declared; they are passed through, gate by gate, and the humanoid economy's confusion is that its members are at five different gates describing themselves identically. This report grades every claim in the category against a single ladder of evidence, from the stage to the balance sheet.

EVIDENCE REQUIRED RISES LEFT TO RIGHT



The Five Gates - what a robot platform must pass between demonstration and industry, and what counts as evidence at each.

Gate one, IT WORKS: the machine performs the task under controlled conditions. Evidence: a demonstration - which the desk's standing manual treats as a claim, interrogated for continuous takes, third-party settings, and intervention counts. Gate two, IT WORKS UNATTENDED: the machine performs without a hidden human, at a disclosed intervention rate. Evidence: autonomy statistics with the teleoperation share on the record - the gate over which the industry has just split into open doctrinal camps. Gate three, IT SELLS: someone who is not a laboratory pays for it. Evidence: shipments with buyer mix attached; the audited filing outranks the press release by the desk's standing hierarchy. Gate four, IT EARNS: the machine, deployed, produces more value than it consumes, durably. Evidence: utilisation, reliability, attendant ratios, renewal - the numbers this report will show almost nobody publishes. Gate five, IT SCALES: production, supply chain and margin survive volume. Evidence: capacity actually used, component sourcing, gross margin at ascending volume - the filing's territory again.

The gates do the sorting the marketing refuses to. A half-marathon in fifty minutes is a gate-one event of genuine engineering significance and zero gate-three content. An audited five thousand shipments is a gate-three fact that says nothing about gate four. A seventy-five-thousand-unit capacity commitment is a gate-five wager, disclosed. The chapters of Part Two take the gates in order, and the composite verdict at the end is a position on the ladder, not a mood.

The year the demos ended

Three developments made 2025-26 the period the humanoid question became checkable, and each is a genre shift in the evidence rather than a leap in the machines. The first is the arrival of the audited filing: the industry's volume leader submitted a 363-page prospectus in March, putting shipments, prices, margins, buyer mix and capacity under securities law for the first time at scale, three years after the category's first listing - the subject of this desk's filings manual and the spine of this report's order-book chapter. Marketing can now be marked against a legal document, which is a different sport. The second is the arrival of paid, countable work: robots-as-a-service contracts at named logistics and manufacturing sites converted the demo reel into an invoice stream, and one operator's tote counter became the industry's first odometer. The third is the arrival of the rulebook: the international industrial-robot safety standards were revised and brought into force in April 2025, their US adoption published in October, and the first walking-robot-specific standard entered development - which means the category's claims are now being tested simultaneously by auditors, customers and standards engineers, three professions with no interest in choreography.

What did not change is the physics of the pitch. The same season produced a robot half-marathon won in fifty minutes and twenty-six seconds, a boxing league, and a state target of one hundred thousand deployed humanoids by 2027 backed by five ministries - each a real event, each a gate-one or gate-zero artefact, each reported somewhere as evidence of industry. The gap between the marathon and the invoice is the gap this report is built to measure.

PART TWO

The Gates



Gates one and two: autonomy, and the split

Begin with the industry's oldest open secret, because 2025-26 is the period it stopped being one. Behind a large share of the category's most fluid demonstrations sits a human operator - the practice this desk documented as standard, undisclosed industry method, detectable by the questions the demo manual asks: single take? third-party site? who is holding the controller? What changed is that the secret became a strategy war fought in public, and the two poles could not be more usefully opposed.

At one pole, the home-robot pioneer made teleoperation the product. The first consumer humanoid shipped to houses this spring - twenty thousand dollars or a monthly subscription, produced at a newly opened California factory - with a launch design its maker describes plainly: when the robot does not know a task, a company operator takes over remotely, completes it through the robot's cameras, and the demonstration trains the model. Industry estimates put launch autonomy at sixty to seventy percent; the company's own framing is that owners schedule an operator for chores the machine has not learned. The honesty is genuinely to its credit - the reveal that launch demonstrations were operator-driven came within hours of pre-orders opening, and the company confirmed the architecture rather than disputing the observation - but the design decision stands: the first humanoid in the home is, by construction, part machine and part remote worker, with the privacy consequences of a camera-bearing, microphone-bearing, remotely drivable body in a private house consented to in the terms. At the other pole, the most heavily capitalised Western maker planted its doctrine in public: no teleoperation, ever, in the home - eat the autonomy pain now rather than import the privacy and labour questions later. Between the poles sits the deployed fleet's quieter truth: the most commercially proven humanoid in the world operates behind physical barriers, cannot yet detect humans, and works a narrow, verified task loop - autonomy achieved by constraining the world rather than mastering it, which is exactly how industrial robotics has always actually worked.

The technical bet for closing gate two is uniform across camps even as the disclosure ethics diverge: vision-language-action foundation models - the platform vendor's robot-foundation stack, the leading maker's in-house model, the research lineage of cross-embodiment training - fed by exactly the demonstration data teleoperation produces. That is the strategic logic under the home-robot design: every remote intervention is a labelled training example from an environment simulation cannot faithfully produce, the sim-to-real gap being this desk's standing beat. The bet may well pay. The report's finding is narrower and firmer: as of July 2026, no maker publishes the one statistic that would let a buyer locate any platform on gate two - the intervention rate, per task class, over time. Autonomy percentages appear in coverage as industry estimates; they appear in no filing, no spec sheet and no customer disclosure this desk can find. The gate is not failed. It is unlit.

VERDICT Autonomy is the industry's contested gate and its least instrumented. One camp ships teleoperation as an honest product feature; another forswears it as doctrine; the deployed leader sidesteps it with cages and constraint. Until intervention rates are published the way shipments now are, every autonomy claim is a gate-one artefact. What would move it: any maker putting intervention statistics in a filing.

Gate three: the order book

The order book is where the humanoid economy divides into two industries wearing one name, and the honest map requires holding both in view. Measured in units, the industry is Chinese: the volume leader's audited filing records roughly 5,500 humanoids shipped in 2025 - at prices falling from about \$85,000 to about \$25,000 average with gross margin near sixty percent, and a sixteen-thousand-dollar list price on its smallest model - with a stated target of ten to twenty thousand units in 2026 and committed capacity of seventy-five thousand a year; a second Chinese maker ships in the low thousands; and the country's structural position runs deeper than any vendor, holding a large majority of category unveilings since 2022 and, by one bank's supply-chain accounting, roughly seventy percent of the component base. Measured in verified productive work, the industry is Western and startlingly small. The deployment ledger's hardest entries: one logistics robot past one hundred thousand totes at a named warehouse under a paying robots-as-a-service contract running since mid-2024, the same platform through more than a year of eight-hour daily factory shifts at a named bearings plant, a seven-robot fleet at a named automaker's Canadian logistics operation under a February agreement, and an eleven-month automotive pilot that ended with numbers most of the industry has never published for anything: ninety thousand components moved, roughly 1,250 operating hours, 1.2 million steps, better than 99% placement accuracy, in support of thirty thousand vehicles - after which the customer noted, accurately and deflatingly, that no robots remained permanently stationed. Independent tracking of the deployment leader estimates its entire installed base at roughly seventy-five units worldwide, within a wide band, the company never having disclosed a count.

The buyer mix is the gate's honest dial, and the filing set it: more than seventy percent of the volume leader's 2025 humanoid shipments went to research and education buyers. The industrial thesis becomes real, as this desk's coverage has held all year, in the quarter that share drops below half - and no maker East or West is there. What fills the space between the audited thousands and the verified dozens is the announcement economy: a May agreement between a British startup and a German industrial supplier described across the trade press as among the largest disclosed humanoid deals ever, a national airline testing units at its hub airport, a home-market pre-order book, and a scoreboard of memoranda that this report's evidence grades sort below the tote counter. The desk's rule for this chapter is the hierarchy from its manuals, extended one rung: an audited filing beats a customer-counted deployment, which beats a named pilot with numbers, which beats a named pilot, which beats a signed agreement, which beats a demonstration - and the industry's centre of gravity in July 2026 sits between the third and fourth rungs, moving up.

VERDICT The order book is real and split: China owns the units, the West owns the
TRACKING evidence-per-unit, and laboratories still take the majority of what ships. The gate opens
- for anyone - in the quarter a maker's research share falls below half. Every filing from
here on answers this in audited numbers, which is why the desk reads them so you
don't have to.

CHAPTER 06

Gate four: the utilisation void

Here is the gate at which the report's method matters most, because the finding is an absence. An industry is proven at gate four - the machine, deployed, durably producing more value than it consumes - and the numbers that prove it have names: utilisation hours per month, mean time between failures, attendant ratio, service cost per operating hour, and the state of the fleet at month eighteen, after the pilot glow and before the refresh cycle. As of July 2026, this desk can find no humanoid maker, East or West, listed or private, that publishes any of them. The nearest public artefacts are fragments from

customers rather than vendors: one logistics operator's testing framework scoring humanoids at seventy to eighty-five percent of human pick speed, the tote and shift counters of the deployment leader's two mature sites, and the automotive pilot's operating-hour disclosure - roughly 1,250 hours across eleven months, which a reader with a calendar will note averages under four hours a day, a number that says pilot in exactly the way the accompanying press said production.

Into the void flows promotional arithmetic, and the desk's unit-economics work supplies the corrective frame this report adopts wholesale. The circulating math divides a robot's price by a five-year, twenty-hour-day life and announces a wage - a dollar or two an hour - that no deployed machine has ever demonstrated, because the calculation assumes the three numbers nobody publishes: near-total utilisation, negligible attendance, and durability past warranty. The honest equation prices the robot-hour as purchase or lease cost plus integration plus the human attendant fraction plus service, divided by hours actually worked - and on the only public inputs available, the fraction of a human attendant alone can dominate the wage. The industry's own commercial structure quietly concedes the point: the robots-as-a-service model that now carries most Western deployments keeps the asset, the maintenance and the risk on the vendor's books, which is rational precisely when the vendor cannot yet prove to a buyer what gate four requires - and which makes the sector's true reliability data exist, privately, inside the vendors' service ledgers, unpublished. The first maker to open that ledger converts the entire category's credibility problem into a competitive weapon, which is why this report's marker table watches for it by name.

VERDICT EARLY

Gate four is where the humanoid industry would be proven and where it is least documented: no maker publishes utilisation, reliability or attendant ratios, and the promotional per-hour wage math assumes all three. The RaaS boom is the tell - vendors keeping the asset means vendors carrying the unknowns. First utilisation disclosure wins the category's trust, and this desk's attention.

CHAPTER 07

Gate five: capacity, components and the wager

The fifth gate is the industry's loudest, because capacity is a press release while shipment is an audit. The disclosed wagers, in one place: the volume leader's committed capacity of seventy-five thousand humanoids a year against fifty-five hundred shipped - a fourteen-fold bet, on the record; the American flagship's dedicated plant announced at twelve thousand units a year of nameplate capacity; the deployment leader's factory publicly described as running well below its ceiling; and the most-watched program of all beginning third-generation production at its California plant this very month - initial run on the order of a thousand units, a larger Texas line dated 2027, a twenty-to-thirty-thousand-dollar price target - against a claims ledger this desk has audited all year, in which production-readiness was announced for 2023, thousands-in-use for 2024, and ten thousand built for 2025, none of which the record supports. Slippage of roughly a year per year is that program's audited cadence; the July line-start is the first entry with contemporary evidence behind it, and the marker table holds its feet to the initial-run number.

Beneath the assembly wagers sits the component question, and it is the chapter's structural finding: by the supply-chain accounting this report can source, roughly seventy percent of the humanoid component base - the actuators, reducers, sensors and batteries that are the machine - concentrates in China, the same country that took a large majority of category unveilings since 2022, published a five-ministry deployment target of one hundred thousand humanoids by 2027, and opened production lines that publish their process as specification: twenty-four assembly stages, seventy-seven inspection

checkpoints, forty-one simulated work-condition tests per unit, the desk's standing example of numbers offered as hostages to verification. The Western response is explicit and funded - the best-capitalised US challenger raised on a stated anti-China thesis, and sovereign-supply language now appears in European deployment announcements - but as of this report, a Western humanoid at volume is substantially a Chinese bill of materials with final assembly elsewhere, and the price gap says so: sixteen thousand dollars against Western platforms at fifty thousand and up. The falling-price-with-fat-margin signature in the volume leader's filing - eighty-five to twenty-five thousand dollars with margins near sixty percent - is what vertical integration looks like in audited form, and it is the single strongest gate-five evidence any maker has produced.

VERDICT TRACKING

Capacity is a disclosed wager running ten to twenty times shipments across the industry, and the component base concentrates in one country to a degree the price sheets confirm. The gate's audited evidence - falling prices with rising margins - currently belongs to exactly one maker. Watch initial-run numbers against announced nameplates, starting this month in Fremont.

CHAPTER 08

The rulebook arrives

An industry, as opposed to a product category, has rules, and the humanoid economy acquired most of a rulebook during the report period - with one load-bearing page still blank. The international industrial-robot safety standards were revised and brought into force on the first of April 2025, absorbing the separate collaborative-robot specification and, in a shift of philosophy the category should read carefully, replacing the certified robot with the certified application: safety now attaches to the machine, the task, the tooling and the environment together, which is precisely the right frame for a general-purpose machine whose entire pitch is task mobility - and precisely the expensive one, since every new workload is a new assessment. The US national adoption published in October; the European machinery regime hardens further in January 2027; and enforcement in the American deployments runs through the general-duty route, where consensus standards define what a recognised hazard is - meaning a vendor's "it's safe" and a regulator's "show the assessment" are now different sentences with different paperwork. The deployment leader holds the category's first field-inspection pass from a nationally recognised testing laboratory, a genuine first this report records as such.

The blank page is the machine itself. The standards in force were written for robots that are bolted down or wheeled; the first international standard specifically for dynamically stable robots - machines that balance, and therefore machines that fall - is still in development as this report goes to press. Every walking humanoid in every factory today is deployed into that gap, managed by fall-zone calculations, cages and the constraint engineering of the autonomy chapter; every humanoid in a home sits further out still, in service-robot and consumer territory where the operative rulebook is, for now, the privacy regime governing the cameras and microphones of the teleoperation model as much as any safety text. None of this is scandal - standards trail machines everywhere - but the sequencing deserves the ledger: the industry's shipping decade has begun before its walking-robot standard has a publication date, and the first serious fall in a mixed workspace will be adjudicated under rules written for a different body.

VERDICT
EARLY

The industrial rulebook modernised on schedule and the category earned its first real certifications - but the standard for the defining property of a humanoid, dynamic stability, is still a draft while the machines ship. The gap is being managed by cages and consent forms. What would move it: ISO 25785-1 published, and the first platform certified against it.


PART THREE

The Economics



What a robot costs, and what a robot-hour costs

The category's price architecture, assembled from the record: sixteen thousand dollars for the Chinese volume leader's small platform, twenty thousand for the consumer home robot - or a \$499 monthly subscription, the clearest signal yet that even consumer humanoids will be sold as service - a sub-fifty-thousand target published by the best-funded US industrial challenger, a twenty-to-thirty-thousand-dollar target for the highest-profile program, Western industrial platforms today at fifty thousand and up with the research flagship above one hundred and fifty, and a category-wide unit-cost decline of roughly forty percent across 2022-24 by adoption-tracker accounting. Around the hardware sits the actual market structure: robots-as-a-service for industry, subscription for the home, both of which convert the gate-four unknowns into vendor-side risk and buyer-side operating expense - and both of which mean the sticker prices above are decreasingly what anyone pays.



The forecast economy deserves its own paragraph, logged rather than endorsed. One global bank projects two million workplace humanoids by 2035 and three hundred million by 2050; others publish trillion-dollar total-addressable-market curves; and this report treats every such figure by the house rule for projections - as claims about models, not about robots, useful for reading investor behaviour and useless for verifying anything. The verifiable economics live one level down, in numbers this report has already graded: a sixty-percent gross margin on falling prices in one audited filing, a twenty-two-to-twenty-eight-percent first-year labour-cost reduction claimed for well-matched pilots by deployment intermediaries, a seventy-to-eighty-five-percent-of-human pick rate from one customer's testing, and the unpublished service ledgers of gate four. The distance between the trillion-dollar curves and the seventy-five verified robots of the deployment tracker is not a contradiction; it is a duration. The market is pricing the endpoint. This report prices the evidence.

The verifier's toolkit

The reader's field kit for every humanoid claim of the next twelve months, compressed from the desk's manuals into one page. For a demonstration: is it a single continuous take, on a third-party site, and who

is holding the controller - and if the answer to the last is undisclosed, the demo is a claim about choreography. For an autonomy statement: ask for the intervention rate per task class; no rate, no gate two. For a deployment: grade the evidence on the six-rung hierarchy - audited filing, customer-counted work, named pilot with numbers, named pilot, signed agreement, demonstration - and notice that press coverage routinely reports rung five in the language of rung two. For a filing: the desk's six lines - buyer mix first, price curve with margin, revenue mix, capacity versus shipped, related parties, risk factors read as confession. For a capacity announcement: divide by current shipments and file the quotient as the size of the wager. For a per-hour wage claim: ask for utilisation, attendant ratio and service cost, the three inputs the arithmetic assumes and the industry does not publish. And for this document: the fingerprint on the integrity page recomputes from the canonical text with one line of standard tooling, recipe in Appendix B - the same standard we have just spent a report applying to everyone else.

PART FOUR

The Scoreboard



Three years on the record

THE RECORD, DATED

DATE	EVENT	WHY IT CARRIES WEIGHT
2023	The category's first stock-exchange listing puts a humanoid maker under securities law	The audited era begins
Jun 2024	A logistics operator signs the first commercial robots-as-a-service humanoid deployment	The invoice era begins
Late 2024	An automotive pilot concludes with part-counts, hours and accuracy published; a bearings maker takes an equity stake in its robot supplier and puts the fleet on daily shifts	Deployment evidence reaches its highest grade to date
1 Apr 2025	Revised international industrial-robot safety standards come into force, certifying applications rather than robots	The rulebook modernises
Oct 2025	US national adoption of the revised standards publishes; consumer home-robot pre-orders open at \$20,000 - and the launch demos are identified as teleoperated within hours	The rules and the open secret go public in the same month
Early 2025-2026	Factory shifts accumulate: 100,000 totes at one site; eight-hour days at another; a seven-unit automaker fleet signs in February	The tote counter becomes the industry's odometer
Mar 2026	The volume leader files a 363-page prospectus: 5,500 shipped, >70% research buyers, ~60% margins, 75,000-unit capacity	The industry's most honest document to date
Apr 2026	A robot wins a half-marathon in 50:26; independent tracking estimates the deployment leader's global fleet at ~75 units	The gap between spectacle and installed base, quantified in one month
May 2026	The consumer factory opens and home shipments begin; a British-German industrial agreement is reported as the category's largest disclosed deal	The home era and the announcement economy arrive together
Jul 2026	Third-generation production begins at the most-watched program's California plant, initial run ~1,000 units - as this report goes to press	The claims ledger meets an assembly line

Twelve months of checkable markers

The markers this publication will score in public, each phrased so that success and failure are both recognisable. The Fremont number: third-generation units actually produced by year end against the ~1,000 initial run and the program's own history of year-per-year slippage. The volume leader's 2026 shipments against its stated ten-to-twenty-thousand target - answerable from its next filing, which is the point of filings. The buyer-mix dial: any maker's research-and-education share crossing below fifty percent, the quarter the industrial thesis gets its first real data point. The first utilisation disclosure: any maker publishing hours, reliability or attendant ratios - the gate-four unlock this report has argued is the

category's cheapest available credibility. The intervention rate: any autonomy claim backed by per-task intervention statistics in a filing or spec sheet. The walking-robot standard: ISO 25785-1 reaching publication, and the first platform certified against it. The home-robot autonomy curve: the consumer platform's teleoperation share at year end against its sixty-to-seventy-percent launch estimate, by any independent measurement. And the nameplate audit: announced factory capacities against any disclosed output, ours to compile if no one else does. Where each lands, the next edition prints it beside what we predicted.

CHAPTER 13

What would change our verdict

The composite verdict below moves toward VERIFIED on evidence, not on volume: a maker crossing the buyer-mix line with utilisation numbers attached; intervention rates published and independently checked; the walking-robot standard in force with certified platforms; audited shipments meeting audited capacity at sustained margins. It moves toward CONTESTED on the patterns this desk already tracks: another year of capacity announcements outrunning shipments by an order of magnitude; the announcement economy continuing to out-publish the deployment ledger; autonomy statistics remaining unlit while demonstrations escalate; or a serious safety event adjudicated in the standards gap this report has described. And it would move immediately, in either direction, on the one document this category owes the record and has never produced: a maintenance log, opened.

CHAPTER 14

The verdicts

The claim under test was that general-purpose robots have crossed from demonstration to industry. Against the Five Gates: gate one is VERIFIED and always was - the machines work, spectacularly, under conditions their makers control. Gate two is CONTESTED and unlit - the teleoperation split is now doctrine versus product feature, and no one publishes the rate. Gate three is TRACKING - audited thousands in China, verified dozens in the West, laboratories still the majority buyer. Gate four is EARLY and empty - the industry's proof lives in service ledgers nobody opens. Gate five is TRACKING as a disclosed wager - fourteen-fold capacity bets on the record, a component base concentrated in one country, and exactly one maker with audited scale economics. The rulebook is EARLY, modern for arms and a draft for walkers.

The composite is TRACKING, and the sentence-length version of this report is this: the humanoid economy in July 2026 is a real industry on the supply side - factories, filings, standards, service contracts - attached to a demand side still measured in the hundreds of verifiably working robots, doing narrow tasks, superbly counted where they are counted at all. The crossing from demonstration to industry is genuinely underway, and it is happening at the pace of the tote counter, not the keynote. The machines have learned to walk. The industry is still learning to show its work - and this desk will be reading the logs, the filings and the fall zones when it does.

The claims ledger

Every material claim in this report, with the status The Verifier assigns it and the source that carries it. Statuses: VERIFIED (checked against the primary record), TRACKING (credible, single-sourced or vendor-reported, monitored), CONTESTED (sources genuinely conflict, or the claim is an absence we assert), EARLY (a category-state judgement).

VERIFIED	Unitree's March 2026 Shanghai prospectus (363 pages) discloses ~5,500 humanoids shipped in 2025, >70% to research and education buyers, prices falling from ~\$85,000 to ~\$25,000 average with gross margin near 60%, a \$16,000 list price on the G1, and committed capacity of 75,000 humanoids a year	Unitree SSE filing via Rest of World analysis (Mar 2026); The Verifier filings coverage
VERIFIED	UBTech became the category's first listed humanoid maker in 2023	HKEX listing record (2023)
TRACKING	Unitree targets 10,000-20,000 humanoid shipments in 2026 - company statement, checkable against its next filing	Company statements; sector trackers (2026)
VERIFIED	Agility Robotics' Digit surpassed 100,000 totes moved at GXO's Flowery Branch facility under a robots-as-a-service agreement signed June 2024 - the category's most verified productive work	GXO and Agility disclosures; deployment tracking (2024-26)
VERIFIED	Digit has run ~8-hour daily shifts at Schaeffler's Cheraw, South Carolina plant since early 2025; Schaeffler took an equity stake in November 2024; Toyota Motor Manufacturing Canada signed a RaaS agreement in February 2026 with seven or more units active on RAV4 logistics; Mercado Libre signed in December 2025	Company announcements; deployment tracking (2024-26)
TRACKING	Independent tracking estimates Digit's global installed base at roughly 75 units (plausible range 40-150); Agility has never disclosed a count, and its factory is described as running well below capacity	New Market Pitch deployment tracker (Apr 2026)
VERIFIED	Agility holds the category's first OSHA-recognised NRTL field inspection pass	Agility announcement; tracker corroboration
VERIFIED	Figure's 11-month pilot at BMW Spartanburg logged ~90,000 components moved, ~1,250 operating hours, ~1.2 million steps and >99% placement accuracy in support of ~30,000 X3 vehicles; BMW stated no robots remained permanently stationed as of late 2025	Figure and BMW disclosures; pilot coverage (2025-26)
VERIFIED	BMW expanded its humanoid programme to Plant Leipzig in February 2026 with Hexagon's wheeled AEON for EV battery assembly, with full-scale testing targeted for summer 2026, and established a Centre of Competence for Physical AI	BMW announcements (Feb 2026)
VERIFIED	Apptronik's Apollo is deployed in pilots with Mercedes-Benz (component transport up to ~25kg and quality checks) and works with Jabil and GXO; the company partners with Google DeepMind on foundation models and designs for field-replaceable modular joints and swappable ~4-hour battery packs	Apptronik and customer announcements (2025-26)
TRACKING	Apptronik's cumulative funding is reported at \$935 million with a \$5.5 billion valuation and an explicitly stated aim of beating Chinese competitors; Figure's valuation is reported at \$39 billion as of September 2025 - secondary-sourced figures	Sector reporting (2025-26)
VERIFIED	Tesla began converting Fremont capacity after ending S/X lines in May 2026 and starts third-generation Optimus production in the July-August 2026 window at an initial run on the order of 1,000 units, with a Texas line dated 2027 and a \$20,000-30,000 price target; current internal deployment is primarily for data collection	Production coverage (May-Jul 2026); The Verifier Optimus claims ledger

VERIFIED	The Optimus programme's public claims - production-ready 2023, thousands in use 2024, 10,000 built 2025 - are unsupported by the record, with slippage of roughly a year per year	The Verifier, "The Optimus Ledger" (Jul 2026) and cited primary statements
VERIFIED	1X's NEO consumer humanoid opened pre-orders in late October 2025 at \$20,000 or \$499/month, began shipping from a newly opened Hayward, California factory on 1 May 2026, and uses an Expert Mode in which company teleoperators complete tasks the robot cannot, with owner consent, as an explicit design and training strategy	1X announcements; launch coverage (Oct 2025-May 2026)
TRACKING	Industry estimates place NEO's launch autonomy at 60-70%, with roughly a third of encountered tasks requiring human intervention - estimates, not company-published statistics	Launch-period reporting (2025-26)
VERIFIED	NEO's launch demonstrations were identified as teleoperated within hours of pre-orders opening, including in Wall Street Journal testing, and the company confirmed the human-in-the-loop architecture	WSJ coverage (Oct 2025); subsequent reporting
VERIFIED	Figure's leadership has publicly committed to not using teleoperation in home deployment, citing privacy and the choice to solve autonomy directly	Public statements by Figure Leadership (2025-26)
VERIFIED	The most-deployed commercial humanoid operates behind physical barriers, cannot yet detect humans, and cannot complete a full shift on one charge, per documented deployment constraints	Deployment tracking and site reporting (2025-26)
VERIFIED	No humanoid maker publishes utilisation hours, mean time between failures, attendant ratios or intervention rates as of July 2026 - an absence this report asserts and invites correction on	Survey of maker disclosures, filings and spec sheets through Jul 2026
VERIFIED	GXO's testing framework has scored humanoids at 70-85% of human pick speed	GXO pilot disclosures
TRACKING	Deployment intermediaries claim 22-28% first-year labour-cost reduction for well-matched pilot tasks - vendor-side figures	Integration-sector materials (2026)
VERIFIED	Category unit costs fell roughly 40% across 2022-24, with current platform prices spanning ~\$16,000 to over \$150,000	Adoption trackers and price lists (2025-26)
TRACKING	China accounts for a large majority of humanoid unveilings since 2022 and roughly 70% of the component supply chain, per Morgan Stanley's accounting	Morgan Stanley humanoid research (2025)
VERIFIED	China's MIIT and five ministries target 100,000 deployed humanoids by 2027; a Guangdong line opened in March 2026 publishing a 24-stage, 77-checkpoint, 41-test process specification	Policy documents; line-opening coverage; The Verifier robotics coverage
VERIFIED	A humanoid won a Beijing half-marathon on 19 April 2026 in 50:26 - a capability demonstration this report grades as gate-one evidence only	Event coverage (Apr 2026)
VERIFIED	ISO 10218-1:2025 and 10218-2:2025 came into force on 1 April 2025, absorbing ISO/TS 15066 and shifting certification from robots to applications; ANSI/A3 R15.06-2025 published in October 2025; the EU Machinery Regulation 2023/1230 applies from January 2027	Standards records; compliance analyses (2025-26)
VERIFIED	ISO 25785-1, the first international standard specifically for dynamically stable (walking) robots, remains in development as of this report - no published walking-robot safety standard exists	Standards-development records (2026)
VERIFIED	US enforcement for humanoid deployments runs through OSHA's General Duty Clause with consensus standards defining recognised hazards - there is no humanoid-specific OSHA standard	Regulatory analyses (2026)
VERIFIED	A May 2026 agreement between UK startup Humanoid and Schaeffler was reported across trade press as among the largest disclosed humanoid deployment deals; Japan Airlines is testing humanoids at Haneda - agreements and pilots, graded below counted work on this report's evidence hierarchy	Deal and pilot coverage (May 2026)

TRACKING	UBS projects 2 million workplace humanoids by 2035 and 300 million by 2050 - projections, logged as claims about models	UBS research via sector coverage
VERIFIED	Robots-as-a-service and subscription models carry most Western deployments and the consumer launch, keeping assets, maintenance and reliability risk on vendor books	Contract structures across GX0, Toyota, Schaeffler-Humanoid and 1X offerings

Method and fingerprint

The fingerprint on the integrity page is computed by the house recipe, identical across this series and every article this publication ships. Take the canonical text file published alongside this report (report-03-canonical.txt) - the report's manuscript of record, front matter through Appendix C, structural markers included. Decode any character entities, remove every whitespace character of any kind, convert the result to lower case, and compute the SHA-256 digest of the UTF-8 bytes; the hexadecimal digest must match the fingerprint printed in this document exactly. Any correction produces a new canonical text, a new fingerprint, and a dated entry in the corrections log on the report's page. The fingerprint proves the text you hold is the text we published; the ledger, the sources and the markers are for whether it is true.

Sources

Filings and primary: Unitree Robotics SSE prospectus (March 2026), via Rest of World's prospectus analysis; UBTech HKEX listing record (2023); ISO 10218-1:2025 and 10218-2:2025; ANSI/A3 R15.06-2025; ISO 25785-1 development record; EU Machinery Regulation 2023/1230; MIIT humanoid deployment guidance. Deployments and operators: GXO and Agility Robotics disclosures (RaaS agreement, tote milestones); Schaeffler announcements and equity disclosure; Toyota Motor Manufacturing Canada agreement (Feb 2026); Mercado Libre agreement (Dec 2025); Figure and BMW Spartanburg pilot disclosures and BMW's Leipzig/Hexagon AEON expansion (Feb 2026); Aptronik, Mercedes-Benz and Jabil announcements; Tesla production coverage (May-Jul 2026); 1X Technologies NEO launch materials and Hayward factory opening (May 2026); Japan Airlines Haneda pilots; Humanoid-Schaeffler agreement coverage (May 2026). Tracking and analysis: New Market Pitch Digit deployment tracker and factory-adoption analyses (2026); Morgan Stanley humanoid supply-chain research (2025); UBS workplace-humanoid projections; GXO pilot KPI disclosures; Wall Street Journal NEO testing (Oct 2025); The Robot Report on teleoperation strategy; sector deployment guides (2026). The Verifier's own reporting drawn on throughout: "China's Humanoid Ledger," "The Optimus Ledger," "How to Read a Robotics Filing," "How to Read a Robot Demo," "Teleoperation: the Industry's Open Secret," "The Unit Economics of a Robot That Works," "Sim-to-Real," "Why Humanoids Are the Hardest Way to Build a Robot," and "Industrial Policy With Legs."