

WHITE PAPER

From pilot to productivity

What it takes to scale humanoid robots in industrial operations



Executive positioning

Labor shortages, increasing operational complexity and rising productivity pressure are exposing the limits of traditional automation. At the same time, humanoid robotics is moving out of the laboratory and into industrial pilot deployments. The critical question is whether humanoid robots can deliver reliable, economical and repeatable productivity at scale.

Organizations that establish a structured deployment strategy today will be better positioned to capture value as the technology matures.

Executive Summary

Industrial humanoid robotics is entering a productivity phase. The relevant management question is shifting from whether humanoid robots can demonstrate isolated tasks to whether they can deliver reliable output within real industrial operations, over productive periods, with acceptable integration effort and a credible business case.

Current evidence indicates that warehouse and manufacturing applications are the most realistic near-term commercialization domains. Industrial pilots and use-case assessments concentrate on material movement, carry handling, warehouse loading and unloading, line-side logistics, selected assembly support and inspection-related activities. These use cases are attractive because they combine labor intensity, repeatability and structured operating environments while still requiring flexibility that conventional automation may not always provide.


The central scaling barrier is operational. Runtime, uptime, reliability, maintenance effort, process integration, safety governance and workforce adoption determine whether a pilot becomes a productivity system. Energy strategy is particularly consequential: battery capacity, charging concept incl. swapping approach and infrastructure footprint influence staffing models, shift design and total cost of ownership.

This white paper proposes a structured readiness lens for industrial decision-makers. Organizations should evaluate humanoid robotics through four questions: which processes have the strongest operational fit, which technical constraints limit productive uptime, which systems and workflows must be integrated, and which organizational capabilities are needed to operate and improve the deployment over time.

ADVANCED ROBOTICS

Humanoids can be deployed in a variety of use cases in different industries

HUMANOID ROBOT USE CASES



In this White Paper

After reading this whitepaper, decision-makers will be able to:

- Identify which industrial use cases are most likely to scale successfully
- Evaluate humanoid deployments against productivity rather than demonstration metrics
- Understand the operational barriers that determine deployment success
- Assess energy, integration and workforce requirements before major investments
- Apply a structured readiness framework to prioritize pilot opportunities

Introduction

The cost of inaction in industrial automation is becoming more visible. Warehouses and manufacturing plants face labor availability constraints, higher flexibility requirements and pressure to improve productivity without redesigning every process around fixed automation. Humanoid robots are relevant because they are intended to operate in environments designed for people, rather than requiring a fully redesigned production or logistics system. This creates an opportunity, but it also raises a stricter standard for proof: a humanoid pilot must eventually translate into measurable productive output.

The first phase of humanoid robotics was shaped by demonstrations. Robots walked, grasped, sorted items and interacted with people. The next phase is being shaped by industrial evidence. Operations leaders now need to know whether humanoid robots can meet task-specific cycle times, complete

enough productive hours per shift, integrate with existing operational systems and remain safe and reliable in environments that contain people, equipment and process variability.

This whitepaper focuses on the gap between technical feasibility and operational productivity. It uses internal research on use cases, OEM strategies, battery and charging systems, and E/E-software infrastructure to identify the factors that separate a promising pilot from a scalable deployment.

Methodology

The analysis synthesizes four evidence streams that are required to evaluate industrial humanoid viability. Each stream addresses a different layer of the deployment question: where humanoids can create value, which platforms and suppliers are active, which energy constraints shape uptime, and which software and integration requirements influence scaling.

Evidence stream	Analytical contribution	Primary decision question
Industrial use-case analysis	Assessment of warehouse and automotive manufacturing tasks, including cycle time, reach, payload, runtime and process complexity.	Where can humanoids create near-term operational value?
OEM landscape analysis	Benchmarking of leading humanoid developers, product maturity, pilot tasks, partnerships and business models.	Which suppliers and capabilities are relevant for industrial deployment?
Battery and charging analysis	Evaluation of battery architecture, charging strategy, energy efficiency, safety, mass, cost and reliability trade-offs.	Can the robot operate productively within the required shift and infrastructure model?
E/E, software and compute analysis	Review of onboard compute, control paradigms, deployment software, data collection, inference and fleet management concepts.	How can robots be deployed, monitored, integrated and improved over time?

The evidence base includes expert interviews, desktop research, public information, patents and internal assessment frameworks. Where evidence is qualitative, this whitepaper treats the finding as directional rather than definitive. Where numerical requirements are available, the analysis preserves the original ranges and qualifiers.

Analysis & Findings

1. Industrial use cases have become the near-term commercialization path

Industrial environments provide the most credible near-term path for humanoid scaling because they combine structured workflows with remaining manual work. Warehouse and manufacturing environments are more predictable than household or broad service environments, yet they still contain tasks that

require flexibility, human-like reach, object handling and adaptation. This makes them suitable for the early phase of general-purpose robotic deployment.

Warehouse applications are particularly relevant where the work is repetitive, labor intensive and difficult to fully automate with fixed equipment. Use-case analysis indicates that humanoids have the greatest potential in tasks that require flexibility and real-time adaptation while maintaining repetitive execution. This is especially relevant in semi-automated logistics environments where conventional automation may deliver high throughput but can lack flexibility or require significant implementation effort in brownfield sites.

Automotive manufacturing shows a more selective opportunity profile. High-volume automotive production is already highly automated in areas such as metal processing, body-in-white and painting, but manual work remains meaningful in assembly, final inspection, inbound logistics and material handling. The implication is that humanoid deployments should not start with a generic plant-wide ambition. They should start with task-level screening of repeatability, complexity, safety, accuracy and decision-making requirements.



2. Productivity requirements are already quantifiable

The transition from pilot to productivity requires measurable task requirements. In warehouse operations, high-potential static use cases require approximately 550 picks or moves per hour, while dynamic use cases require approximately 300 picks or moves per hour. Typical warehouse reach requirements are around 0.5–1.0 m horizontally and up to 2.0–2.2 m vertically, with payload requirements up to 25 kg and runtime expectations between approximately 5 and 8 hours depending on the task type.


Automotive manufacturing requirements are similarly specific. Dynamic tasks such as lifting, positioning and assembling components typically require 60–90 picks or moves per hour per component in medium- and high-potential use cases. Semi-static tasks with limited travel requirements can require approximately 240 picks or moves per hour per component. Many automotive use cases target a payload up to 15 kg, horizontal reach of 0.5–1.0 m, vertical reach of 1.5–1.8 m and minimum runtime aligned with an 8-hour shift.

These figures create a practical implication for decision-makers: humanoids should be evaluated against operational KPIs rather than demonstration milestones. A robot that can complete a task once may still be unfit for deployment if it cannot meet the required cycle time, task accuracy, runtime, safety and handover requirements across a full operational period.

Requirement area	Warehouse reference range	Automotive manufacturing reference range	Operational implication
Throughput / cycle time	Static: approx. 550 picks or moves per hour; dynamic: approx. 300 picks or moves per hour.	Semi-static: approx. 15 sec/component; dynamic: approx. 40–60 sec/component.	Pilot KPIs need to be measured against process-specific productivity requirements.
Reach	Approx. 0.5–1.0 m horizontal; up to 2.0–2.2 m vertical.	Approx. 0.5–1.0 m horizontal; 1.5–1.8 m vertical.	Human-designed workspaces define the practical geometry of deployment.
Payload	Up to 25 kg for many relevant tasks.	Up to 15 kg in the analyzed value-chain requirements.	Payload claims must be evaluated in relation to task frequency, handling safety and endurance.
Runtime	Approx. 5 h for transportation tasks and 8 h for stationary tasks.	Minimum 8 h as one-shift reference.	Energy strategy directly influences staffing, charging and fleet-size assumptions.

The humanoid potential is highest for value chain steps with a low automation degree today and relatively low picking rates

HUMANOID POTENTIAL MAP



For each value chain step, the **automation level of the current solution must be evaluated to determine humanoid robot potential**

- Humanoids are designed to be general purpose, bringing **flexibility in brownfield warehouses**
- Compared to fully-automated systems, humanoids can **scale with demand** during seasonal fluctuations
- Humanoids have the potential to **free up humans from repetitive, dangerous tasks**
- As humanoids' **AI & learning models improve**, their ability to **adapt to a multitude of SKUs** will increase

3. Current pilots show that operational constraints still dominate scaling

Pilot activity is concentrated in tasks such as line-side material distribution, warehouse handling and sorting, tote consolidation, palletizing, trolley transportation, screw box transportation, wheel-related handling, sheet-metal placement, quality inspection and pick-and-place activities. These tasks are operationally meaningful because they sit close to existing labor-intensive workflows and are easier to observe, measure and improve than highly variable household or service activities.

However, pilot evidence also shows the distance still to travel. Reported industrial pilots have required combinations of teleoperation, autonomy, supplier support, battery improvement and task-specific workflow definition. In one pilot context, ten robots per factory and approximately fifty supporting staff were reported, with runtime of around 3–3.5 hours regularly achieved. In another context, battery runtime and charging options were significant constraints shaping whether the robot could cover a meaningful portion of a shift.

The implication is that early deployments must be designed as learning systems. A pilot should not only test whether a robot can execute the task. It should measure productive uptime, interventions, charging or swapping effort, maintenance load, safety stops, process disruption, data requirements and whether the same behavior transfers across product variants or operational conditions.

4. Energy management is a deployment architecture decision

Battery systems affect more than robot runtime. Energy architecture influences robot mass, balance, packaging, task capability, charging time, maintenance labor, infrastructure complexity and safety. The core trade-off is that increasing energy content can improve runtime, but additional mass may degrade balance and energy efficiency and can create integration constraints.



The analysis of humanoid battery systems highlights several interdependent factors: runtime between charges, nominal uptime, lifetime uptime, maintenance accessibility, cell aging, charging capability, human involvement and infrastructure footprint at the end-user site. Assessing these variables in isolation can distort deployment economics. For example, a larger battery pack may increase energy capacity but also add weight, which raises energy consumption.

Charging concepts are also not interchangeable. Manual plug-in, self-recharging, battery swapping, wired operation and wireless charging each imply different levels of downtime, human involvement, infrastructure complexity and operational robustness. For industrial fleets, standardization and modularization may become increasingly relevant, especially if swapping or automated charging concepts mature.

Energy decision	Main deployment effect	Key risk if ignored
Battery capacity	Defines runtime potential but affects mass, packaging and balance.	A larger battery can reduce charging frequency while degrading efficiency or integration quality.
Swappable vs. fixed pack	Shapes serviceability, downtime and charging-room logistics.	A concept that looks attractive in the robot may be impractical in shift operations.
Charging infrastructure	Determines charging downtime, footprint and plant integration.	Insufficient charging design can force additional fleet size or manual support.
Thermal management and BMS	Influences safety, reliability, aging and usable operating window.	Poor control of aging or safety constraints can undermine the TCO case.

5. Software and integration determine whether robots become operational assets

Humanoid deployment is increasingly a software and systems-integration challenge. Robots need task assignment, diagnostics, monitoring, telemetry and integration into production or logistics systems. In industrial settings, the robot is not only deployed physically; its workflow must be defined, validated, updated and connected to the systems that coordinate the work environment.

Internal analysis indicates that deployment models focus on workflow definition and enterprise integration, as well as data collection and AI policy improvement. Examples include connections to the warehouse management system or manufacturing execution system, as well as the development of fleet-management and mission-control capabilities.

Data quality also matters for productivity. In one logistics use-case example, more curated data reportedly produced stronger throughput improvement than a larger non-curated data set. This indicates that the quality and relevance of deployment data may be more important than volume alone. For executives, the implication is that future productivity improvement may depend on software integration, data governance and continuous learning, not only on mechanical capability.

6. Scaling requires workforce and operating-model readiness

Workforce adoption is a necessary part of scaling. Humanoid deployments can introduce new activities such as robot supervision, battery handling, troubleshooting, safety management, task reconfiguration and data collection. In early pilots, support from robot suppliers and operator training can be substantial. As pilots scale, these activities must become part of the operating model rather than temporary project work.

The operating model should specify who owns task performance, who manages robot readiness, how issues are escalated, which KPIs are tracked, how safety stops are handled and how process improvements are fed back into deployment. Without this structure, a technically promising pilot may remain dependent on exceptional support and fail to become a sustainable productivity system.

This is why the readiness question should extend beyond the robot. Facilities need process readiness, infrastructure readiness, and workforce readiness. The strongest candidates for scale will be processes where these elements can be aligned without disproportionate redesign effort.

Which Use Cases Are Most Likely to Scale First?

Near-term scaling potential is highest where the task is repetitive, structured, labor intensive but still requires enough flexibility to make humanoid morphology useful. The analysis points toward logistics and material-handling applications before complex inspection or highly dexterous assembly tasks.

Use case family	Near-term scale potential	Reasoning	Main constraint to validate
Material transportation and line-side logistics	High	Structured movement of materials is labor intensive and appears repeatedly in pilot tasks.	Runtime, fleet coordination and interaction with existing material-flow systems.
Tote handling, loading and unloading	High	Tasks are repetitive and directly linked to warehouse labor demand.	Throughput, grasp robustness, object variability and charging downtime.
Palletizing, stacking and buffer-zone handling	Medium to high	Repetitive handling in structured spaces can create a measurable productivity case.	Payload, reach, accuracy and interaction with conveyors, AGVs or lifting devices.
Basic assembly support	Medium	Selected tasks may fit humanoids when they are repeatable and not too complex.	Dexterity, task accuracy, cycle time and product-variant complexity.
Complex final inspection	Lower near term	Inspection can require nuanced judgment and high accuracy under time constraints.	Perception reliability, data requirements and risk of missed defects.



Implications for Industrial Decision-Makers

The findings imply that humanoid robotics should be evaluated as an operational transformation topic rather than as a robot procurement decision. The most important decision is which use case, integration model and operating concept can produce measurable output under real conditions.

Industrial pilot activity is accelerating across manufacturing, logistics and warehouse operations. While technology maturity remains uneven, organizations that establish evaluation frameworks, pilot capabilities and operational experience early will be better positioned to capture value as deployment readiness improves. The immediate opportunity building the capabilities required to identify scalable use cases and generate operational learning today.

A structured readiness assessment can reduce risk by separating feasibility from productivity. Feasibility asks whether the robot can complete a task. Productivity asks whether the robot can complete the task with the required cycle time, uptime, quality, safety, integration effort and support model. This distinction helps executives decide where to pilot, when to scale and where additional maturity is still required.

FEV Consulting is well positioned to support this transition because the challenge spans operations strategy, manufacturing and logistics process expertise, electrification and battery systems, software integration, supplier assessment and business-case development. A structured readiness assessment helps organizations identify high-potential use cases, quantify operational requirements, evaluate deployment constraints and establish clear decision criteria for scaling.

Recommended executive questions

1. Which tasks combine labor intensity, repeatability and enough flexibility to justify humanoid deployment?
2. What productive uptime is required for a credible business case, and how will charging or swapping be integrated?
3. Which systems must the robots connect to before they can become operational assets?
4. Which workforce roles, safety procedures and support processes are required for scale?

Key takeaways and next steps

Industrial humanoid robotics is entering a new phase of industrial adoption. Future deployment success will depend on four factors:

- The right use case
- A viable runtime and energy strategy
- Robust process and software integration
- Workforce and operating-model readiness

Industrial pilot activity is accelerating across manufacturing, logistics and warehouse operations. Organizations that build evaluation frameworks and gain operational experience early will be better positioned to identify scalable opportunities and capture value.

From Insight to Action

Before investing in larger-scale deployments, organizations should address three questions:

1. Which processes offer the strongest deployment potential?
2. What productivity levels are required to support a viable business case?
3. Which technical, operational and organizational barriers must be addressed?

A structured readiness assessment can help:

- Prioritize high-potential use cases
- Quantify productivity requirements
- Evaluate deployment constraints
- Define a roadmap for industrial adoption

The outcome is a fact-based view of where humanoid robotics can create measurable operational value and how to move from pilot projects toward scalable deployment

Appendix

A. Humanoid Readiness Assessment Checklist

Dimension	Screening question
Process fit	Does the task combine repetition, labor intensity and moderate variability?
Operational KPI fit	Are cycle time, payload, reach, accuracy and runtime requirements defined?
Energy readiness	Can the process accommodate charging, swapping or docking without undermining productivity?
Integration readiness	Which WMS, WES, MES, MOM, ERP or fleet-management interfaces are required?
Safety readiness	Which human-robot interaction conditions, emergency stops and operating boundaries are required?
Workforce readiness	Who supervises, maintains, resets and improves the deployment after the supplier team leaves?
Business-case readiness	Which cost, uptime, productivity and quality metrics determine scale-up approval?

B. Suggested Pilot Scorecard

Metric	Purpose
Productive uptime	Measures how much time the robot performs value-adding work.
Cycle time / throughput	Compares robot output with process requirements.
Intervention rate	Tracks how often human support is needed.
Charging or swapping downtime	Quantifies energy-related productivity loss.
Maintenance effort	Captures routine support and unscheduled service needs.
Safety events	Tracks emergency stops, slowdowns, virtual-cage events or other safety constraints.
Transferability	Measures whether the behavior works across product variants, stations or shifts.

C. Key Terms

Humanoid robot: A general-purpose robotic platform designed to operate in environments originally built for human workers.

Productive uptime: The share of available time in which the robot performs productive work rather than charging, waiting, being reset or undergoing maintenance.

Battery swapping: A deployment concept in which depleted robot batteries are replaced with charged packs to reduce downtime.

Autonomous docking: A charging concept in which the robot navigates to a charging interface and connects without manual plug-in.

Fleet management: Software used to monitor, assign, coordinate and optimize multiple robots in an operating environment.

Operational validation: A pilot stage that tests performance against real process KPIs, rather than only technical task execution.

Contacts



Dr. Hagen Wegner

Director

FEV Consulting

wegner_h@fev.com



Benjamin Knobloch

Director

FEV Consulting

knobloch@fev.com



Dr. Dominik Boemer, PMP

Senior Manager

FEV Consulting

boemer@fev.com