INDUSTRY REPORT

IEX 2026 / DBR77

Physical AI in American Manufacturing

What can work, what must be proved, and what to do next



DBR77 INDUSTRIAL INTELLIGENCE

DBR77 Robotics Inc.

dr Piotr Wiśniewski · Justyna Laskowska
Paweł Mroczkowski · Torian Richardson · Bartek Straszak

AI-generated conceptual illustration

OCTOBER 2026

DBR77 RESEARCH / THE TEAM BEHIND THE REPORT

The team behind the report

Published by DBR77 Robotics Inc.

dr Piotr Wiśniewski · Justyna Laskowska · Paweł Mroczkowski · Torian Richardson · Bartek Straszak

dr Piotr Wiśniewski is the founder and Global CEO of DBR77, a Harvard Business School General Management Program alumnus and former CEO of Maflow. His experience connects manufacturing leadership, operational transformation and industrial technology.

Justyna Laskowska is Head of U.S. Market Development at DBR77. A former Director of the Centre for Academic Entrepreneurship and Technology Transfer at Nicolaus Copernicus University, she brings extensive experience in strategic partnerships, innovation management and technology commercialization.

Paweł Mroczkowski is Executive Director Western Europe at DBR77. An entrepreneur with experience in Poland and Germany, he holds an MBA from Hult International Business School and completed executive programs at Harvard Business School and Yale School of Management.

Torian Richardson is CEO of DBR77 USA and a former Global Director at NVIDIA. His background spans operations, workforce transformation and technology. He is a Harvard Business School General Management Program alumnus.

Bartek Straszak is Chief Operating Officer at DBR77 and a former McKinsey & Company professional. His experience combines operational transformation, digital technologies and manufacturing process optimization, with a focus on applying innovation to industrial operations.

Educational and career affiliations identify the contributors' experience and do not imply institutional endorsement.

AI-assisted research and editorial production

AI tools supported research, source analysis, drafting, analytical graphics, editorial design and typesetting. Generated illustrations are identified where used. DBR77 Robotics Inc. retains editorial responsibility.

DBR77 INDUSTRIAL INTELLIGENCE

CORPORATE AUTHOR

DBR77 Robotics Inc.

DBR77 RESEARCH / IMPORTANT NOTICES & DISCLOSURES

Important Notices & Disclosures

Research reference date: 18 October 2026. Source corrections: 19 October 2026. Design edition: 20 October 2026. Prepared for IEX 2026.

Purpose. Prepared by DBR77 Robotics Inc. for IEX 2026, this report provides general information and analysis of Physical AI in manufacturing. It does not provide site-specific engineering, safety, legal, financial or investment advice, a product certification or an offer of securities.

Evidence. The analysis combines selected public research, official statistics, company disclosures, published interviews and the authors' interpretations. Company accounts are attributed to their publishers. The report is an analytical synthesis, rather than a representative field survey or an independent audit of the organizations discussed. Source definitions, dates and operating conditions govern interpretation. The research reference date is October 18, 2026, with identified source corrections through October 19, 2026.

Models and outlook. Synthetic calculations explain the consequences of declared assumptions. They are not customer outcomes, market prices, valuations or promises of return. The outlook describes conditional pathways. Historical results do not guarantee future performance, and information may change after the reference date.

Industrial use. Deployment decisions require application-specific engineering, risk assessment, validation and compliance work. The report's diagrams and recovery examples are explanatory frameworks, not operating instructions. Nothing here authorizes bypassing protective systems or operating a particular installation. Manufacturers, integrators, employers and operators retain their applicable responsibilities.

Commercial interests. DBR77 develops and markets industrial technology and has a commercial interest in this field. Advertisement pages are identified separately. References to other companies, educational institutions, experts or trademarks do not imply sponsorship, endorsement or agreement with the conclusions. Product descriptions should be evaluated within their stated scope.

AI and images. AI supported research and editorial production, as disclosed on the contributor page. Generated illustrations are labeled and are not documentary photographs of deployments. DBR77 Robotics Inc. retains editorial responsibility.

Third-party rights. Referenced publications, quotations and trademarks remain subject to the rights of their respective owners. Attribution does not grant permission to reproduce or commercially reuse third-party material. This report grants no license to that material.

© 2026 DBR77 Robotics Inc. Third-party material remains the property of its respective rights holders.

DBR77 RESEARCH / EXECUTIVE FINDINGS

Executive findings

Physical AI is becoming an industrial capability, but the market uses the name for several different technologies. For American manufacturers, the opportunity is to make changing work easier to automate and coordinated production easier to manage. The following judgments connect the technical progress with a practical investment agenda.

1. The near-term opportunity is selective adaptation

Perception, spatial reasoning and learned action policies can reduce the engineering needed for variable tasks. The strongest initial applications have a clear operating envelope and a result the plant can measure. The commercially significant advance is less work to adapt, commission and support the next task. Chapters 1–3 explain the capability and machine choices.

2. Established automation is part of the opportunity

Industrial arms, collaborative applications and mobile robots already perform useful work at scale. Learning-based systems can extend that installed base as well as create new applications. Humanoids deserve attention where mobility, reach and reuse across valuable tasks justify their additional complexity. A plant strategy should retain several forms of automation. Chapter 3 separates emerging Physical AI cases from established operating comparators.

3. American competitiveness needs an application-level answer

Reshoring, regional skills, supply-chain resilience and industrial policy affect the setting for investment. They do not replace the demand and cost assumptions of a specific factory. Smaller producers often need accessible engineering and support; high-volume plants need dependable improvements at constrained operations. Chapter 4 connects national context to those different buying problems.

4. Coordination is a credible next control point

As intelligent assets multiply, shared resources and conflicting local objectives become more important. DBR77's thesis is that a common operating context can connect process learning, simulation and human knowledge with better plant decisions. Existing industrial suites, integrators, equipment providers and independent layers can compete to supply it. The opportunity is recurring improvement in order performance and coordination cost. Chapters 5 and 7 develop this argument and its business implications.

5. Demand and exception work determine the return

More technical capacity creates value when there is a profitable use for it. Assistance, recovery, integration and service can change the preferred investment. The report's synthetic comparison shows different demand conditions favoring process improvement, an incumbent extension or a local AI package. Supplier contribution and funding needs require a separate ledger. Chapter 6 makes those mechanisms explicit.

6. Recovery and retained competence determine sustainable scale

A dependable application needs people who can diagnose exceptions, verify the physical state and authorize restart. Cybersecurity, protective functions and change control remain part of its design. Over six months, comparable operating records matter most. Over three years, task and site transfer become central. Over ten years, interoperability, accountability and local competence may shape the market as much as model capability.

DBR77 RESEARCH / CONTENTS

Contents

01	The production question	7
02	From seeing to doing	9
03	Which machine, which task?	17
04	American factory reality	23
05	From decision to operation	27
06	The economics of exceptions	33
07	Who captures value?	39
08	Work and organizational change	43
09	Risks and recovery	46
10	Outlook: the next industrial bottlenecks	51
11	Business decisions: buy the next justified step	53

Reference sections

Technical note: 56. Sources & Notes: 57. About DBR77: 59.

DBR77 RESEARCH / KEY TERMS AND EVIDENCE

Key terms and evidence

The report uses the following operating terms consistently. They describe what an application does and what a buyer can evaluate.

Term	Meaning in this report
Accepted output	Work that meets the operation's stated quality, task and timing requirements.
Operating envelope	The conditions and actions for which an application is designed and authorized.
Intervention / exception	Human assistance or a departure from normal execution that must be included in the operating record.
Recovery	The work required to diagnose an interruption, reconcile physical state and restore authorized operation.
Physical readback	Observation that verifies the actual process result, beyond acknowledgement of a command.
Decision authority	The right to recommend, schedule, command, stop or restart a defined operation.
Transfer	The extent to which capability, engineering and support can be reused on a changed task or another site.

Evidence used

This report synthesizes public research, official statistics, company accounts and published interviews and talks. Interview analysis uses publisher transcripts or published speech text. Company-reported results retain attribution. Synthetic calculations and hypothetical scenarios are identified as analytical examples. The report presents no new representative survey or independent audit of plant returns.

FROM A CLAIM TO A PRODUCTION DECISION

DEFINE THE CLAIM
What action, operating envelope and accepted result?

→

EXAMINE THE EVIDENCE
Which denominator, comparator and recovery record?

→

ASSIGN THE DECISION
Who may act—and what would change that decision?

01 / THE PRODUCTION QUESTION

The production question

A part arrives differently from one cycle to the next. It may be rotated, partly hidden, slightly deformed or mixed with another variant. The production question is how to handle that variation at a cost the customer can justify. Sometimes the answer changes the process. Sometimes it changes what the machine can perceive and decide.

Table 1. Four labels, four different roles

Label	What it means here	What it does not prove
CORE	AI inference materially contributes to selecting an embodied system's physical action.	Safe, autonomous, profitable or production-ready operation.
ADJACENT	Spatial/world models, human-facing assistance and plant operating systems that inform, constrain or coordinate physical work.	Membership in the embodied core or direct control.
BASELINE	Process redesign, conventional automation, local control and human work used as real comparators.	"Old" or inferior; deterministic execution can be the right answer.
ECOSYSTEM	Simulation, compute, integration, procurement and marketplaces enabling deployment.	Intelligence in the purchased machine or a completed outcome.

A fixture might remove the variation. A conventional cell might handle it within a known envelope. A person might remain the best adaptable system. Or a learning-based machine might make valuable variation manageable without redesigning the process around every exception. Physical AI matters when it changes that choice—not when it merely gives an old machine a fashionable label.

The opportunity is to make useful variation cheaper to handle: shorter product runs, changing part presentations and tasks that currently consume substantial engineering attention. The commercial test is the complete application over time, including integration, assistance and return to service. That is where a reusable model can become a production advantage and where a promising demonstration can still fail to become a business.

DBR77's central thesis concerns the next level of that opportunity. As intelligent assets multiply, plants will need a shared operating context connecting process learning, simulation, human operational knowledge and the consequences of decisions. This could be provided by an existing industrial stack, an integrator-led federation or a new operating layer. Its value should appear in better order fulfillment, fewer coordination losses and lower effort when work changes. The architecture can remain distributed while sharing the information that a decision actually needs.

What belongs inside the term?

There is no single boundary shared by the original sources examined for this report. The disagreement

is useful because it reveals three different objects hiding under one name.

In the narrow, embodied meaning, AI participates in selecting the action of a machine. A camera and model perceive a situation; a policy or planner chooses what the robot should do; motors make the choice physical. NVIDIA's institutional glossary uses a wider autonomous-system frame that includes robots, vehicles, cameras and smart spaces, and allows a system to perform or orchestrate physical-world actions.[1]

A second meaning centers on spatial intelligence: representing objects, relationships and change in three-dimensional environments. It may help a robot, but it can also assist a person or generate a simulated world without actuating anything. Fei-Fei Li draws the distinction sharply: "One is about saying things. The other one is about seeing and doing things." In context, she contrasts language models with world models operating across physical and digital environments; her comparison directs attention to the representations needed for action.[2]

A third meaning is architectural. Vision-language-action systems connect images and instructions to motor outputs; embodied-reasoning systems may instead supply spatial interpretation and planning to an existing controller. Google DeepMind's own materials distinguish a VLA that adds physical action as an output from an embodied-reasoning model that can connect to lower-level controllers.[3] This makes "Physical AI" neither a synonym for a humanoid nor a claim that one foundation model should control an entire plant.

The report therefore uses four explicit labels:

This taxonomy is an editorial choice, not a market standard. Borderline cases stay visible. An AI scheduler approved by a dispatcher is ADJACENT unless it participates in an embodied action loop; a fixed-sequence robot is BASELINE unless qualifying AI inference affects its action. A bad or unsafe CORE system remains in the category. Governance, readback and verified results determine maturity; they do not define failure out of existence.

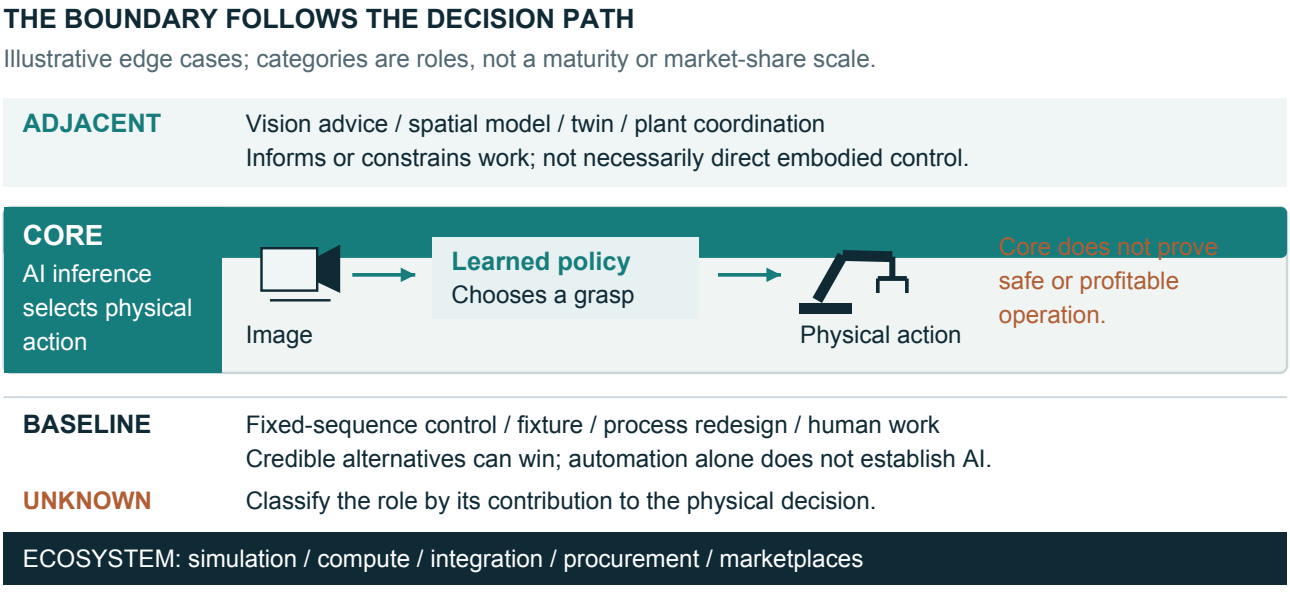
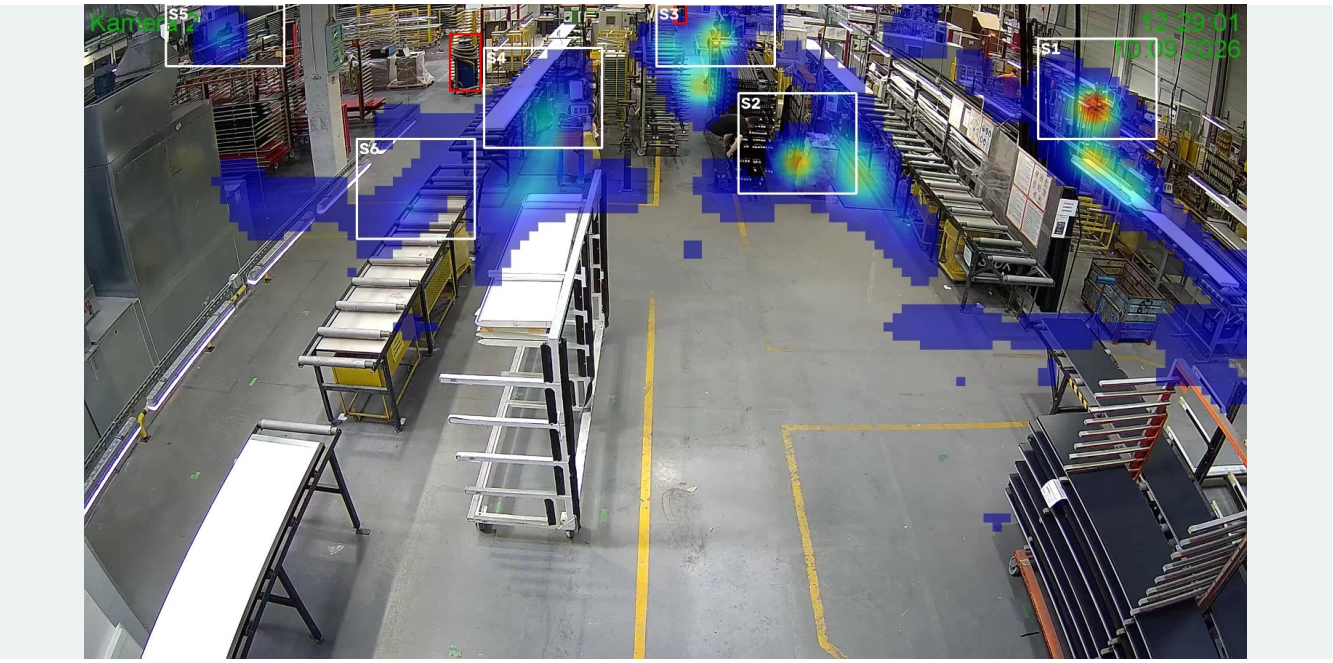


Figure 1. A practical taxonomy of Physical AI. Embodied action, adjacent decision support, conventional comparators and enabling services play different roles. This is the authors' working taxonomy.

02 / FROM SEEING TO DOING

From seeing to doing

What changed in machine intelligence — and what a factory still has to prove



DBR77 camera-analysis illustration; no measured operating improvement is established. The overlay illustrates spatial activity analysis; it does not establish task success or decision authority.

Machine intelligence changes the economics of variation when knowledge learned elsewhere helps solve a new task with less engineering. For a variable part, that capability includes estimating its position, choosing a grasp, adapting the movement and checking the result. The important technical question is which of those abilities transfer when the part, presentation or machine changes.

From visual appearance to task state

Consider our hypothetical part. Naming it is one achievement. Estimating its position, distinguishing the intended surface from an obstruction, and determining whether it can be lifted are different achievements. Knowing that a grasp succeeded is different again. A fluent description of the scene might be impressive while leaving precisely the uncertainty that matters to production unresolved. "Understanding" should therefore describe a tested capability, not an assumed human-like mental state.

For this task, ask what observation would change the next decision. If two parts overlap, is the visible edge sufficient to select a grasp, or should the system seek another observation? If the object is correctly identified but belongs to the wrong order, more visual accuracy will not resolve the mismatch. If a placement signal is missing, should the cell repeat the action, investigate or stop? These are distinct information problems. Treating all of them as computer vision obscures which component must improve and who must verify the result.

In a May 2024 MIT CSAIL talk, roboticist John Leonard distinguished visually attractive scene reconstruction from geometry and prediction useful for robot decisions.[4] That distinction helps specify what a perception system must contribute to the next movement.

The managerial consequence is direct: purchase a capability against a decision requirement, not against the attractiveness of its display. In our example, the specification should say what uncer-

tainty about the part must be resolved, within what time, and what happens when it cannot be resolved. A beautiful three-dimensional reconstruction might help engineering understand the workspace. A simpler estimate might be enough for the next movement. Neither should be judged by the other's purpose.

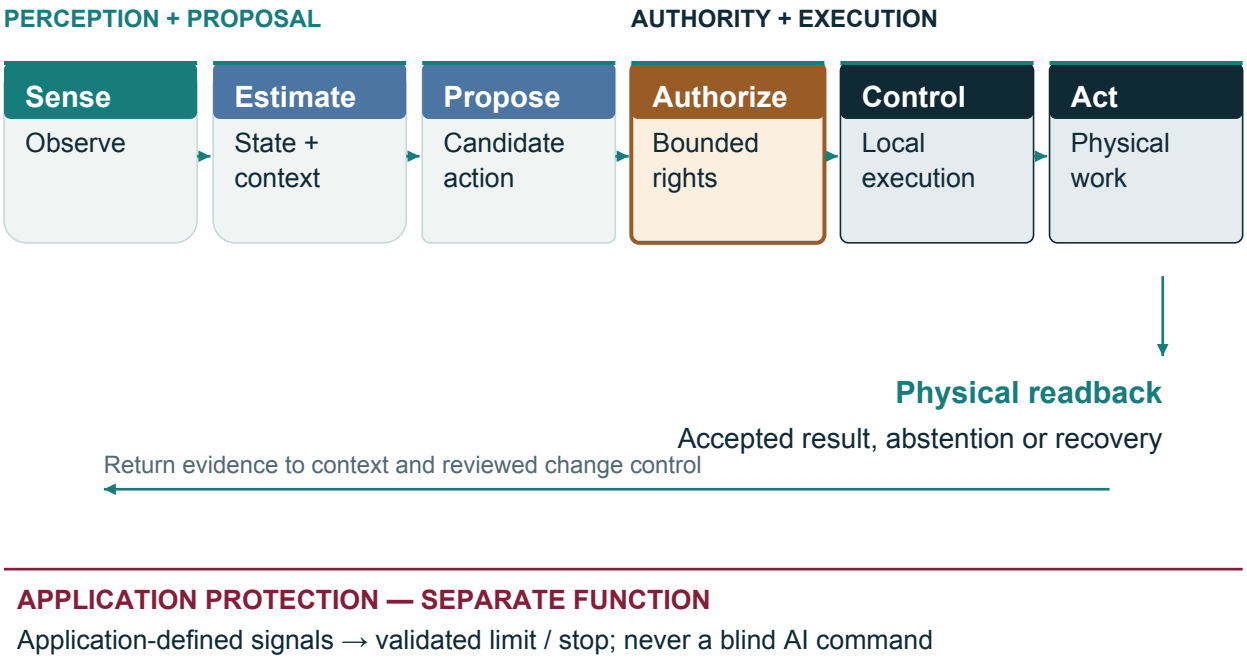


Figure 2. From observation to physical result. Observation, state estimation, planning, authorization, local control and verification contribute distinct functions to the application.

The important shift: connecting representations to action

Google DeepMind's 2023 RT-2 work provides a concrete example of the vision-language-action, or VLA, approach. It combines visual and language pretraining with robot data and represents robot actions as model outputs. Camera images and instructions can consequently participate in selecting physical actions, rather than ending with a written description. Its reported experiments concern research manipulation; they do not establish a factory's cycle time, useful uptime or economics.[5]

The distinction matters because language can become an interface to a task, and prior visual knowledge can become useful beyond a fixed list of hand-coded descriptions. The hypothesis worth testing is that a manufacturer could adapt to a new part presentation or instruction with less bespoke work. That would change the cost of flexibility, not abolish the need to engineer the application. A model's output still has to fit the machine, the process and the permitted operating conditions.

A VLM, or vision-language model, connects image and language representations. A VLA adds an action-producing role.[5]

Those labels do not tell a buyer whether a system chooses a high-level skill, predicts a short movement sequence, or operates behind another planner. Nor do they specify how fast the complete loop responds. The purchasing question is not "Does it have a VLA?" but "Which decision does the model own, and what remains outside it?" A supplier should be able to answer without replacing an architectural explanation with a brand name.

There is no need to make modular engineering and learned policies ideological opponents. In our illustrative task, one design could identify geometry,

construct candidate grasps and pass an explicit plan to a controller. Another could learn more of the observation-to-action mapping. A third could use a learned component only where variation defeats the simpler approach. These are comparison options, not measured rankings. The right test holds the task and acceptance condition constant, then compares adaptation effort, useful output, recoverability and the evidence available when something fails.

Why learning promises more than another programmed motion

Octo, published at Robotics: Science and Systems in 2024, studied a generalist policy pretrained on 800,000 robot trajectories. In its six fine-tuning setups, Table I reports an average success rate of 72%, compared with 20% for training from scratch and 15% for the VC-1 representation baseline. Each setup used approximately 100 target demonstrations and 20 evaluation trials. The paper also reports substantial degradation for novel behaviors and limits its evaluated embodiments to single- and dual-arm manipulators.[6]

This is a meaningful positive result: prior robot experience can provide a better starting point than learning each tested task from scratch. It is not a claim that the model arrives ready for arbitrary work. The distinction between those statements is commercially important. Reuse can be valuable well before general autonomy exists. A buyer should ask whether adaptation requires fewer demonstrations, less engineering, shorter commissioning, or less retraining after a change—and then measure the claimed saving against an equally well-supported conventional alternative.

Learning methods also answer different questions.

Learning from demonstrations asks which behavior to reproduce from examples.[6]

Teleoperation can supply those examples through human-directed robot actions.[7]

Reinforcement learning uses experience to improve sequential decisions against a reward objective. The trade-off includes immediate and later consequences, exploration and generalization.[8]

For the comparison here, ask whether simulation provides an appropriate place to study interaction, without treating every virtual success as a physical result. These are possible development routes, not mandatory stages of one universal pipeline. A purchase specification should identify which route is actually used and where human effort remains.

NVIDIA's December 2025 GR00T N1.6 research report describes added teleoperated data, task-specific post-training and experiments on several embodiments. Its discussion explicitly notes continuing difficulty with language following and out-of-distribution generalization. It also describes sensitivity to the relationship between pretraining and target-task distributions. These qualifications come from the model's developers, alongside their reported improvements.[7]

These results support investment in reusable robot capabilities. The commercially important experiment is whether adaptation to the next permitted task takes fewer demonstrations, fewer engineering hours and less support while preserving quality. Progress in that comparison can create value before general-purpose autonomy is achieved.

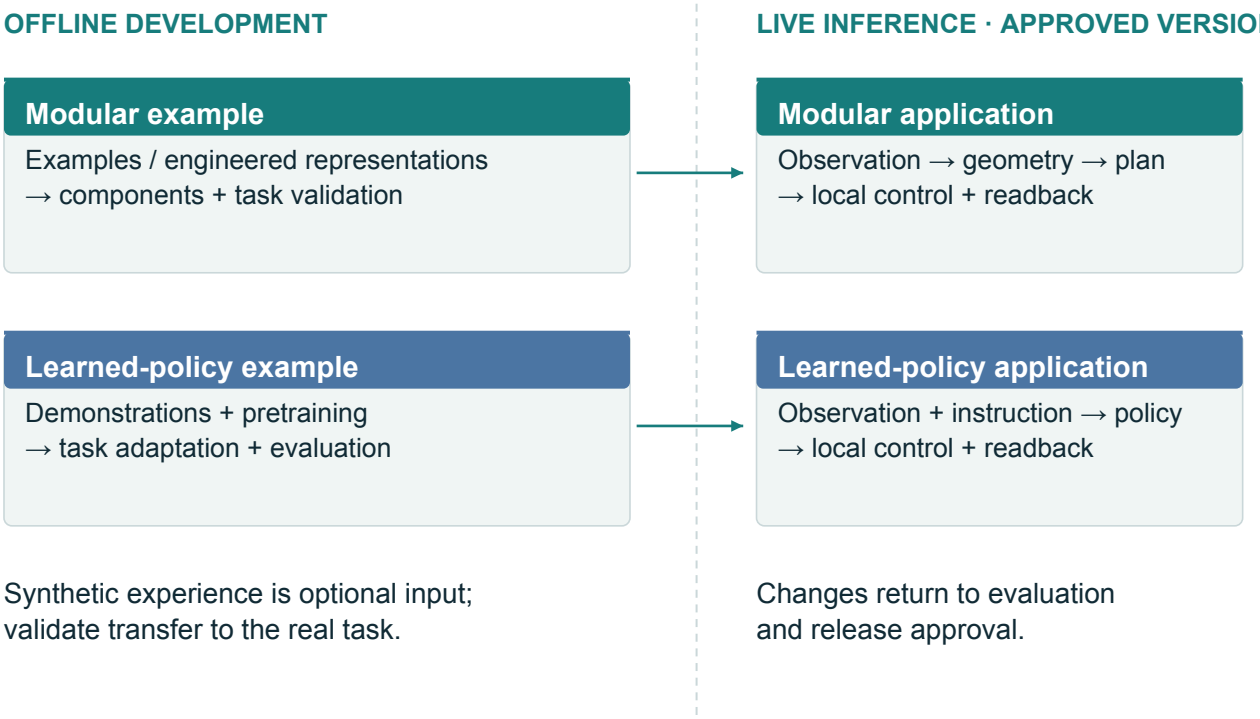


Figure 3. Two illustrative learning routes. Demonstrations, pretraining, adaptation and evaluation can connect through different pipelines. Synthetic experience is an input to be validated, not a substitute for evidence on the real task.

Table 2. Different purposes, different purchasing questions

Purpose in the proposed comparison	Question a buyer must resolve	What would not establish it
Generate an environment	Does variation preserve the features relevant to the test?	Visual realism alone
Predict action consequences	Are the predicted changes accurate within the intended conditions?	A plausible next frame without outcome validation
Support planning in a representation	Does the selected plan improve the accepted task result?	A persuasive verbal explanation
Model an operational process	Is the prediction credible for this decision and current process?	Synchronization or a three-dimensional display alone

Authors' comparison framework, informed by the sources above.[9, 10, 11] Not a standard, exhaustive taxonomy or vendor scorecard.

Four different promises hidden inside “world model”

The phrase “world model” can conceal disagreement about the problem itself. For procurement, this report separates four purposes: generate environments, predict what an action will change, support planning in a compact representation, and model a process for a specified operational decision. These purposes can coexist in one system. They are not a ladder from weak to strong intelligence, and none automatically supplies the others.

DeepMind’s August 2025 Genie 3 announcement describes interactive generated environments, while acknowledging limited direct action space, imperfect geographic fidelity, multi-agent challenges and interaction lasting minutes rather than extended hours. These are properties of the release described at that date, not a statement about current access or every subsequent model.[9]

Such generation could be useful for exploring environmental variation. But our hypothetical plant would still need to know whether the generated variations preserve the relationships relevant to its task. A scene that looks like a workstation does not establish the correct contact behavior of a part, the state of a clamp or the consequence of an intervention. The issue is not whether the image looks convincing. It is whether the model is fit for the decision being tested.

In his March 2024 conversation with Lex Fridman, Yann LeCun advocated prediction and planning in abstract representation space.[10] This approach emphasizes the compact state needed to evaluate actions, offering a different design objective from generating a visually detailed environment.

That disagreement is useful because it redirects the discussion from the word “intelligence” to the computation needed for a decision. Must a system render the next image, predict a compact state, search alternative actions, or estimate a queue? The answer depends on the use. A scheduling model may need credible processing and changeover assumptions without rendering a photorealistic factory. A manipulation model may need a reliable account of geometry and contact without modeling the whole production network.

Shao, Hightower and Schindel connect manufacturing-twin credibility to its intended use, verification, validation and uncertainty assessment throughout its lifecycle.[11] The relevant question is whether the model can support the particular decision being made.

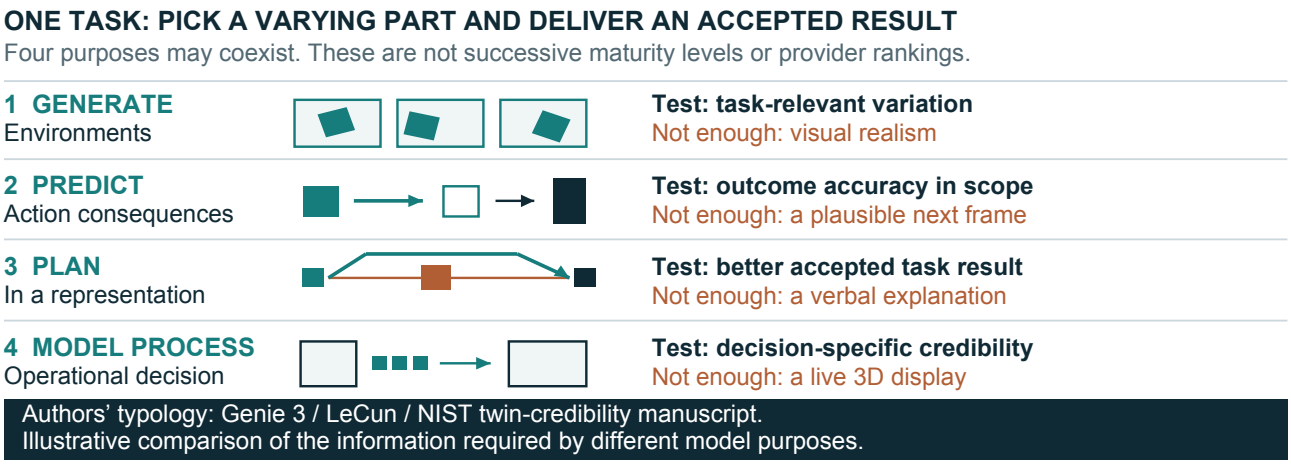


Figure 4. Different world-model purposes. Generating an environment, predicting action consequences, supporting planning in a representation and modeling an operational process require different tests. Visual realism alone does not establish decision credibility.

Generalization has more than one boundary

The 2024 COLOSSEUM benchmark evaluates manipulation under 14 perturbation axes across 20 simulated tasks, including changes in lighting, objects and camera conditions. Its reported degradation demonstrates why performance close to the training setup is not the whole evaluation problem. Those experiments are not a measured failure rate for American factories.[12]

For our part-handling task, distinguish a new position from a new part, a new part from a new behavior, and a new behavior from a new machine. Moving a successful policy between sites adds another question: what changes in tooling, calibration, local process and support? "Works on unseen objects" cannot fill every one of those fields. Generalization becomes commercially meaningful when the supplier declares the dimension of change and the buyer can test its relevance.

The strongest test should include ordinary variation, difficult but permitted conditions, and conditions outside the authorized scope. The last group is not an invitation to execute unsafe actions. It tests whether the system recognizes that the case belongs elsewhere and uses the defined escalation or refusal path. Counting only completed normal cycles can conceal the work shifted to operators. Counting every refusal as a failure can punish appropriate restraint. The task log must distinguish those outcomes before a headline success rate is calculated.

This also changes the role of data. A demonstration collection that covers successful movements may not describe how the application behaves after a dropped part, a changed order or an incomplete restart. For our proposed evaluation, retain the context needed to identify interventions and version changes, not just a curated library of successes. The purpose is not to collect everything indefinitely. It is to make the performance claim reproducible within approved data rights and a declared operating window.

A useful test separates what was learned before installation from what was adapted for this task and what, if anything, changes during operation. The phrase "continual learning" should not collapse those stages. A supplier may collect feedback continuously but release a changed policy only after review; a system may adapt a parameter without retraining its model; another may update its behavior online. These are different propositions to investigate, not capabilities inferred from a product label. Specify which changes are permitted, how the configuration is identified and which evidence must be regenerated after a change.

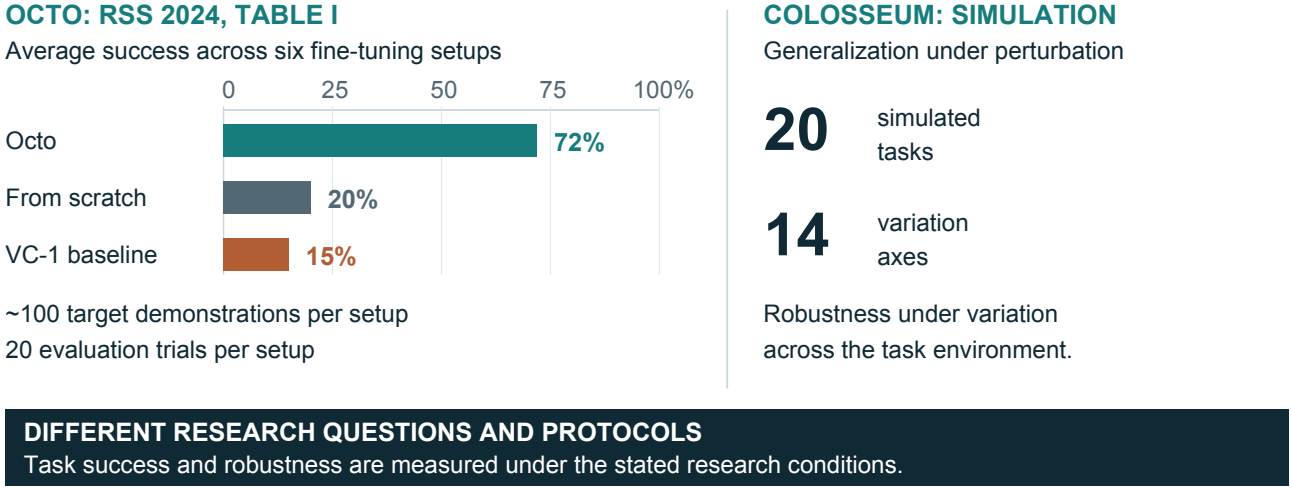


Figure 5. Research on reuse and robustness. Octo compares fine-tuning results across six setups. COLOSSEUM tests variation across simulated manipulation tasks. Their different protocols are shown separately.

From a proposed action to an accepted result

In a June 2025 Siemens podcast, Maximilian Metzner, a robotics practitioner at the Erlangen factory, explained how variable part orientation can lead to different robot motions. The publisher's transcript says, "with AI you also get a certain loss of determinism." His surrounding discussion concerns planning, reach, clearance and brownfield integration—not proof that all AI is stochastic or that conventional automation eliminates uncertainty.[13]

The practical requirement is to state what can vary and what must remain constrained. In the analytical architecture used here, observation informs an operational decision; an authorized command then reaches basic control and actuation; physical readback establishes whether the result was achieved. Protective functions sit alongside that operational chain rather than depending on the same unconstrained inference. This is a reasoning framework, not a certified design or a description of a DBR77 installation.

Return to the missing placement signal. Acknowledgment that a command was received is not evidence that the part was accepted. An instruction to repeat may be wrong if the first placement happened but its signal was lost. A human approval button is not sufficient if the person lacks the state information or time required to intervene. Our evaluation therefore asks who can resolve the uncertainty, what evidence they receive, and what state permits work to resume. The later risk chapter develops these mechanisms; this chapter does not prescribe an installation's safety procedure.

Compute belongs in the same task-level account. Compare total observation-to-decision delay, its variation, connectivity dependencies and behavior when a response is unavailable. A mean response time does not establish that a particular deadline is always met. Edge and remote execution are options to evaluate against these requirements, not automatic indicators of safety or superiority. A model update changes the configuration under test; prior results should not silently transfer to it.

The production test

A September 2026 Siemens/ARC podcast distinguishes established industrial machine learning and vision from newer generative applications, and discusses the difficulty of updating long-lived operational systems without losing confidence in their operation. The conversation supplies a current incumbent perspective, not an independently measured adoption rate or comparative plant outcome.[14]

This distinction prevents two expensive errors: treating every useful industrial AI application as new, and treating every new model as a ready replacement for installed engineering. For the hypothetical task, the next investigation should hold acceptance quality and task boundaries constant, include the best fixture or conventional system, and measure all attempts, interventions, recovery time and changes in conditions. Technical progress matters when it improves that comparison, including the work required to adapt and maintain the system.

The opportunity is larger than a better pick-and-place demonstration. If adaptation becomes cheaper and more reliable, some work currently too variable to automate economically could become accessible. That remains a conditional business thesis, not a measured market forecast. Its test is demanding but understandable: does reusable intelligence reduce the total burden of handling necessary variation while delivering work the factory can accept? That is the bridge from seeing to doing—and from an impressive capability to something worth buying.

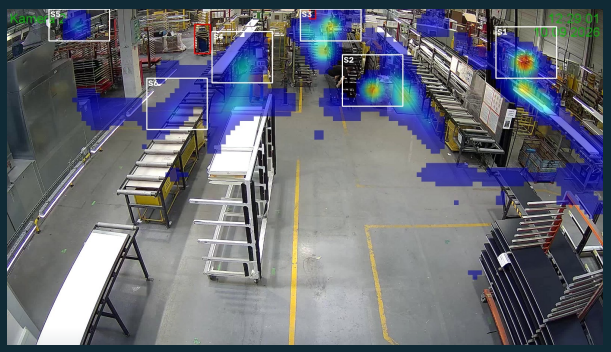
ADVERTISEMENT / DBR77

AI Vision

See what your machines can't tell you.

VISION CAMERA

Selected events.
Structured data.



DBR77 camera-analysis detail

Turn visible production events into process data - even where machines provide no signal.

CYCLE & WAITING

Understand observable process events.

MATERIAL FLOW

Track objects in defined zones.

EDGE PROCESSING

Convert selected video events into data.

AI-generated setting; inset supplied by DBR77. Not a certified safety system.

DBR77
INDUSTRIAL INTELLIGENCE

EXPLORE THE TECHNOLOGY

AI Vision / iot.dbr77.com

03 / WHICH MACHINE, WHICH TASK?

Which machine, which task?

Flexibility is valuable only when it beats the best way to remove the variation



AI-generated conceptual illustration.

TASK ENVELOPE

Variation · reach / payload · rate · human proximity · travel · failure mode

1. **Process redesign / no robot**
   Remove avoidable variation at source
2. **Industrial arm / cell**
   Stable presentation; repeated high-rate work
3. **Collaborative application**
   Bounded shared work; application risk assessment
4. **AGV / AMR**
   Internal movement; traffic, handoffs and rescues
5. **Humanoid / mobile manipulator**
   Mobility + reach + task reuse must justify burden

ONE COMMON COMPARISON

Accepted output · interventions · recovery · task-switch effort · lifecycle cost

Figure 6. Task first, machine second. Process redesign, industrial arms, collaborative applications, mobile transport and humanoids enter the same task comparison. Each must answer for accepted output, interventions, recovery and lifecycle cost.

The choice of machine follows the work: its rate, geometry, variation and expected changes. A fixture, industrial arm, collaborative application, mobile robot or humanoid can each be the economical answer under different conditions. The comparison should include the complete application and the portfolio of tasks it will actually serve.

The most visible machine is rarely the whole application. Tooling, presentation, sensing, guarding, traffic management, software, recovery and human work determine what the factory buys. A general-purpose body may preserve flexibility across several tasks. A specialized system may turn one operation into a faster and more inspectable cycle. The fair contest is complete alternative against complete alternative on the same accepted work—and then across the changes the buyer actually expects.

Table 3. Task-fit comparison

Alternative	Strongest fit	Burden that can erase the advantage	Minimum pilot evidence
Process redesign / no robot	Variation can be removed at source	May leave ergonomic or labor constraint unresolved	New baseline, residual risk, demand and change cost
Industrial arm / dedicated cell	Stable, repeated, high-rate work	Tooling, guarding and changeover for high mix	Accepted output, cycle distribution, availability, recovery and full cell cost
Collaborative application	Bounded task where human value remains nearby	Safe speed/payload limits, nuisance stops, supervision	Mode-specific throughput, interventions, restart burden and application validation
AGV / AMR	Repeated internal transport	Congestion, unstable handoffs, charging and manual rescues	Mission completion, blocked time, rescue rate and degraded behavior
Humanoid / mobile manipulator	Multiple tasks need mobility, reach and reuse of existing space	Generality adds integration, energy, support and validation burden	Useful autonomous time, accepted units, task-switch cost and same-task comparator

Authors' decision framework, not a vendor ranking or safety determination.

Begin with the task envelope

Five questions narrow the choice before a brand enters the conversation. How variable are the object, location and sequence? What reach, payload, precision and rate are required? Must people share the workspace during motion? Does the task require travel between stations? What happens after the normal cycle fails?

An industrial arm is usually strongest when the work is stable, the part can be presented consistently and the output justifies a designed cell. Its apparent rigidity can be an advantage: constraints are visible, cycle logic is inspectable and high-rate motion can occur inside a protected application. A cobot changes the application possibilities by making certain forms of proximity and shared work easier to engineer, but “collaborative” is not a blanket property of every speed, tool, payload and mode. The complete application still needs a risk assessment and validation.

An AGV or AMR addresses movement rather than dexterous work. It can be the right answer when value comes from reliable internal transport, even if it never grasps a component. Routes, pickup geometry, mixed traffic, blocked time, charging and manual rescue become the important variables. A humanoid or mobile manipulator combines reach,

movement and reuse of human-oriented surroundings. That combination is strategically interesting where several changing tasks share an environment. It is wasteful if one stable operation can be solved with a fixture and an arm.

The no-robot option belongs in the same table. A changed container, gravity feed, poka-yoke, improved workstation or revised batch policy may remove the uncertainty at lower lifecycle cost. "No robot" does not mean "do nothing." It means redesign the process before paying a machine to reason about avoidable disorder.

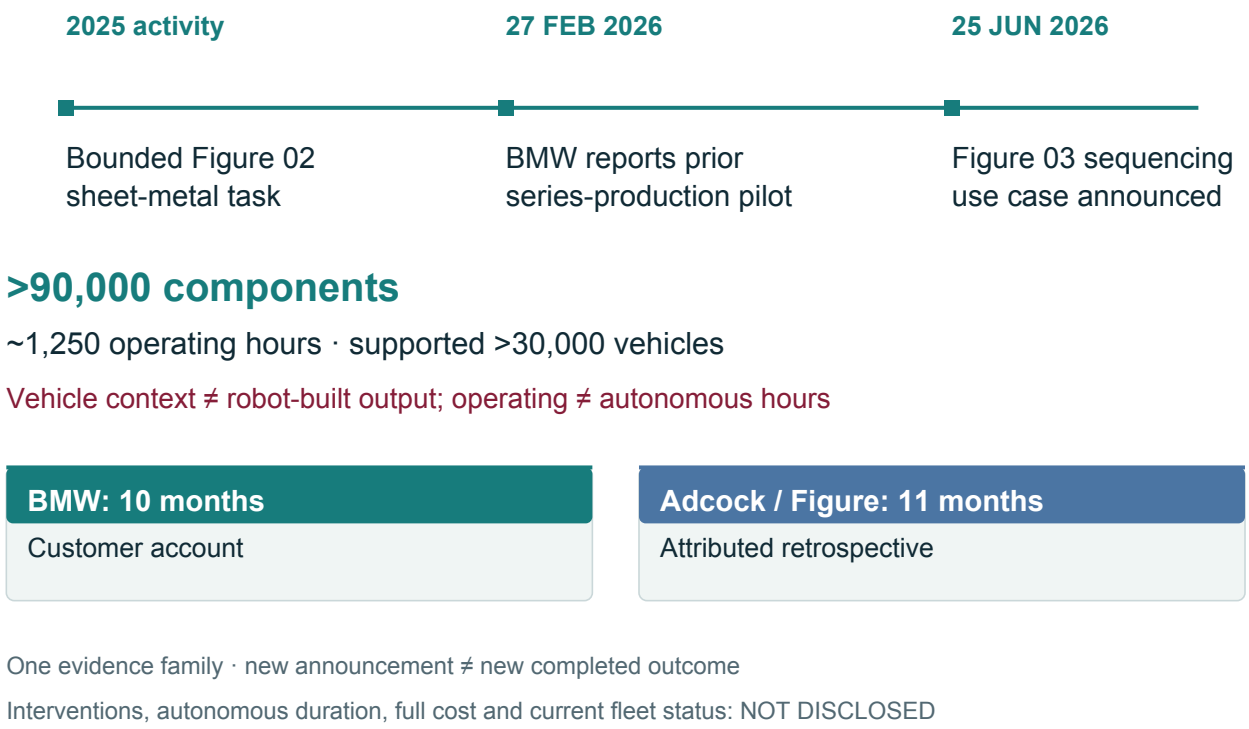


Figure 7. BMW and Figure: operating record and next application. The company-reported 2025 task and June 2026 sequencing project represent different stages of one development relationship.

Emerging Physical AI applications

The emerging humanoid evidence is more substantial than a collection of laboratory films, but it is still uneven. The best records are useful precisely because they expose both progress and missing denominators.

BMW reported that Figure 02 handled more than 90,000 sheet-metal components in about 1,250 operating hours during a ten-month 2025 Spartanburg deployment. The release situated this work within production of more than 30,000 vehicles; it did not say the robot built those vehicles or disclose autonomous duration. In June 2026 BMW announced a Figure 03 logistics-sequencing project. BMW again said ten months; Figure CEO Brett Adcock said eleven.[15, 16]

The Spartanburg record demonstrates sustained work on a defined task in series production. The next commercial questions concern assistance, application cost and transfer to another task. Those fields would turn an operating record into a comparable business case. The June sequencing project extends the development agenda to a different use.

Mercedes-Benz's March 2025 Berlin Digital Factory Campus page described Apollo data collection in production environments, without assigning

all of that work specifically to Berlin. Employees transferred knowledge through teleoperation and augmented reality. The page framed autonomous operation as the next development step and named intralogistics and initial component checks. It supplied no operating-period, intervention, accepted-output or cost denominator.[17]

The Apollo program makes the learning pipeline visible: employees, tools and process knowledge contribute to capability. A buyer can use that example to ask how much data collection repeats at each task and which portion becomes reusable.

GXO's June 2024 investor release said that a proof of concept with Agility Robotics' Digit progressed to a multi-year robotics-as-a-service agreement for a live SPANX warehouse. The release identifies a commercial relationship and an operating setting, but not contracted robot count, useful hours, intervention burden, service-level economics or a conveyor/AMR comparator.[18]

A multi-year service agreement adds a commercial milestone to the technical demonstration. Its relevance is strongest for warehouse and intralogistics handling. Comparing it with machining or assembly requires a separate account of task geometry, quality, cycle requirements and support.

Operating comparators: transport, arms and cobots

The following cases are established automation comparators. They show the operating results a new Physical AI application may need to match. They are not counted as learning-based physical action where the source does not identify such a mechanism.

A 2024 BMW plant release reported nearly 50 tugger trains and more than 140 Smart Transport Robots completing roughly 10,000 trips each workday in Regensburg under cloud traffic control. It did not disclose mission success, blocked time, rescues, network-loss behavior, cost or attribution of the result to the cloud layer.[19]

The example matters because plant-level intelligence may appear first as coordination of specialized machines. If the task is moving known loads between known points, fleet orchestration can be the stronger operating-layer test. A humanoid should not receive credit for flexibility that the use case does not need.

Industrial arms and cobots supply the other baseline. They are not old technology to be dismissed; they are the alternatives a new system must beat.

Universal Robots' RCM Industries story describes two UR10e cells, each tending two dual-spindle lathes. A named RCM director reports an immediate throughput increase of about 15%; the page reports payback in just under a year. It does not state the measurement window, accepted-output denominator, intervention count, downtime, complete cost perimeter or independent financial verification.[20]

FANUC America's Closure Systems International case reports reaching the customer's 97.5% OEE goal and a 25% output increase after six months, described by the supplier as approximately 30 million additional units per month.[21] The account demonstrates why established material-handling automation is a serious comparator. A plant evaluating a similar project would still need its own cost and operating record.

These supplier-hosted cases cannot settle the economics, but they sharpen the comparison. A flexible machine does not compete against a bare robot price. It competes against established cells, integrators and process knowledge that may already deliver high accepted output. Conversely, a dedicated cell should not be credited with flexibility it has not demonstrated when parts, locations or tasks change.

Table 4. What may be concluded from the cases

Case family	Strongest public evidence	Claim that remains out of bounds
BMW / Figure	Bounded series-production task and reported operating activity; later use case announced	Humanoid-built vehicles, ROI, safety certification or multisite superiority
Mercedes / Apollo	Real production environment used for data collection and task development	Autonomous production result or deployment economics
GXO / Digit	Named live-warehouse progression and commercial agreement	Paid productive denominator, manufacturing transfer or low service burden
BMW logistics	Operating specialized transport fleet at plant scale	AI attribution, mission reliability or superior economics
RCM cobots	Repeated machine-tending cells and customer-reported improvement	Independently audited payback or low intervention burden
CSI industrial cells	Six-month company-reported production result	Independent causality, full TCO or safety performance

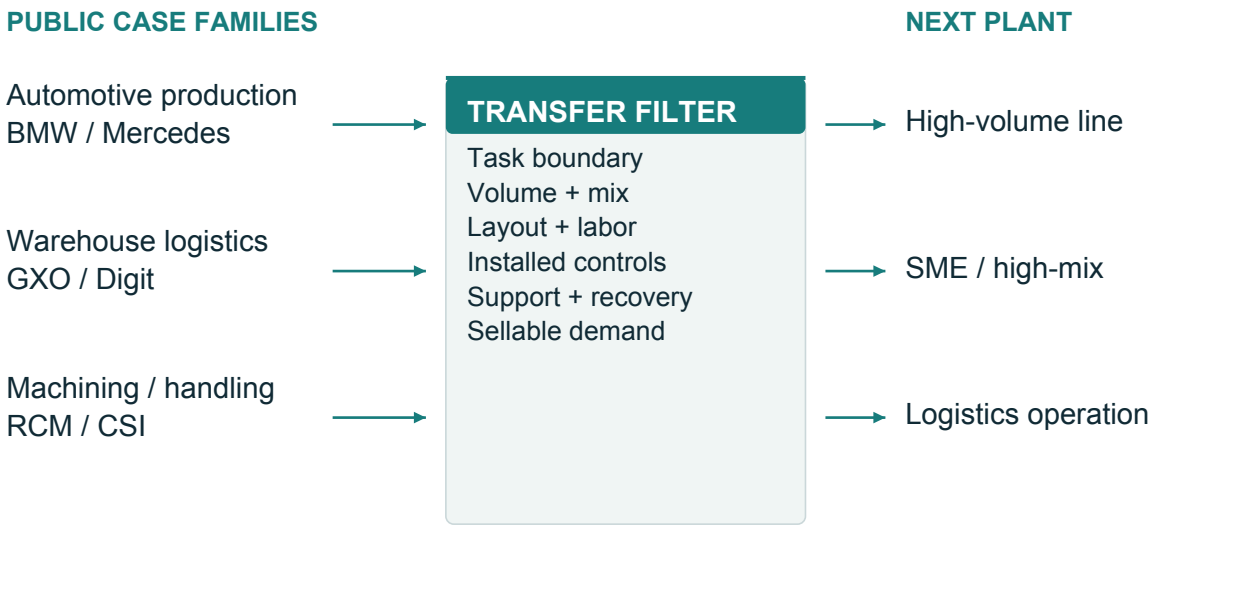


Figure 8. Transferring a result to the next plant. Task, production mix, layout, workforce and support determine whether an operating result is relevant to the next plant.

The price of generality

Generality has at least four sources of value. One machine may perform several tasks. A task may change without rebuilding the cell. A system may operate in space designed for people. A common learning or support layer may improve across deployments. None of those values is automatic. Each has a cost: slower cycle, more complex perception, data collection, supervision, update validation,

energy, maintenance, remote support or a wider failure surface.

Consider an explicitly hypothetical plant with three low-volume jobs that share a work area but not a stable presentation: feeding a machine during one product run, moving containers during changeover, and sorting returned material when the line is idle. A dedicated arm and fixture win the first job's nominal cycle. They may still lose the portfolio decision if each product change requires new tooling, engineering access and lost production while a more general mobile manipulator can switch among all three with less total rework. That does not make the general machine the production winner. It makes it worthy of a bounded pilot whose measured object is task reuse: accepted work across the portfolio, changeover time, interventions, recovery and support.

Now change the portfolio conditions. Demand for the first job becomes stable and high, the other two jobs occur rarely, and a redesigned container removes most presentation variation. The reuse benefit is largely idle. The dedicated arm and fixture regain the advantage; if the container change removes the bottleneck entirely, process redesign or no robot can win instead. The reversal comes from the same task boundary and portfolio, not from a preference for or against a humanoid shape.

The right unit of comparison is therefore a task portfolio over time. A humanoid can lose today's single-task contest and still win if it switches among valuable tasks with low incremental engineering. An arm can dominate if the portfolio is stable and volume is high. A cobot can win where proximity preserves skilled human work. An AMR can win by removing walking and waiting without pretending to solve manipulation. Process redesign can win by deleting the problem.

This creates a disciplined pilot. Define the accepted unit. Record every attempt, refusal, assist, fault and recovery. Separate useful autonomous time from powered-on time. Price tooling, integration, validation, supervision, support and lost production. State the demand condition: capacity has no cash value if there is nothing profitable to produce. Then repeat the test after a meaningful change of part, task or site.

The strongest next evidence would connect operating output to intervention, support and full application cost at a second site. That would establish which gains are repeatable and which depend on the conditions of the first installation.

Choose the next justified machine

The positive case for Physical AI is strong where necessary variation is expensive and reusable perception, planning or manipulation lowers the cost of handling it. The critical case is equally strong: a factory can buy generality it does not use, then pay again for exceptions. The decision is not whether humanoids are real. It is whether a particular embodiment and operating model beat the best alternative for a declared portfolio of work.

04 / AMERICAN FACTORY REALITY

American factory reality

American demand for automation combines an installed robot base, uneven plant readiness and pressure to make domestic production more competitive.

American manufacturing is not one operating environment. A high-volume automotive line, a regional metal fabricator, a contract manufacturer changing products every week and a food plant with seasonal demand can all buy a robot while buying very different economics. The transfer question is therefore not whether a technology has worked "in manufacturing." It is whether the task boundary, variation, volume, skills, installed controls and recovery burden resemble the conditions in which the result was obtained.

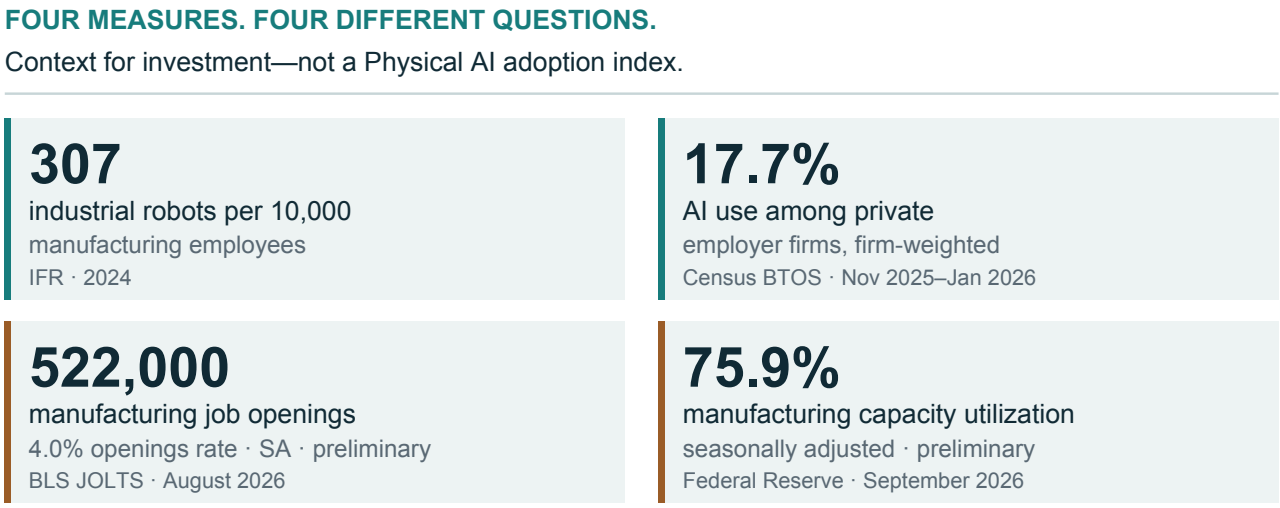


Figure 9. U.S. manufacturing context. Different populations, units and reference periods: IFR robot density in 2024, Census firm-weighted AI use in November 2025–January 2026, BLS manufacturing openings in August 2026 and Federal Reserve utilization in September 2026. Census Table 1 is the basis for AI use; BLS and Fed observations are preliminary.

Four measures that must remain separate

The United States has strong statistics about robots, business use of AI, vacancies and industrial activity. None directly counts factories operating Physical AI.

The International Federation of Robotics reported 307 operational industrial robots per 10,000 manufacturing employees in the United States in 2024.[22] This substantial installed base creates both an adoption route and a demanding comparator for new capabilities. The measure relates robot stock to employment; it does not count factories using Physical AI.

Table 1 of the Census-hosted BTOS working paper reports firm-weighted AI use of 17.7% among private non-agricultural employer firms for the November 2025–January 2026 reference period.[23] This captures self-reported use across business functions in the prior two weeks. It offers context for business diffusion, while plant-floor use requires a more specific measure. The question broadened in November 2025, creating a break with the earlier series.

BLS recorded 522,000 manufacturing job openings, a 4.0% rate, for August 2026; both figures were preliminary and seasonally adjusted.[24] An opening is an unfilled position on the last business day. It does not reveal vacancy duration, occupation, local skill supply or whether automation is the economical remedy.

Federal Reserve G.17 data placed manufacturing capacity utilization at 75.9% in September 2026, below the 78.2% average for 1972–2025; the September observation was preliminary.[25] This is macroeconomic context, not evidence that AI created output, that every plant has spare capacity or that additional automated capacity will find demand.

These measures answer different questions with different units. Combining them into one upward-sloping chart would create a story the data do not contain. Installation flows add context but still count units rather than plants or AI capabilities.

Table 5. What the U.S. measures count

Measure	Counted object	Useful question	Invalid inference
Robot density	Operational industrial robots relative to manufacturing employment	How robot stock relates to the workforce denominator	Share of factories using Physical AI
BTOS AI use	Employer firms self-reporting AI use; Table 1, November 2025–January 2026	Breadth across business functions under the stated questionnaire	Plant-floor or embodied-AI penetration
JOLTS openings	Unfilled positions at month-end	Current hiring context	Structural skills shortage or automation ROI
G.17 utilization	Output relative to estimated capacity	Aggregate manufacturing cycle	AI productivity or site-level available capacity

Reshoring changes the investment calculation

Domestic manufacturing competitiveness depends on more than the cost of a repetitive movement. Proximity to engineering and customers, supply-chain resilience, freight, duties and production quality can all change a location decision. NIST's 2025 U.S. manufacturing economy report, published in February 2026, brings these factors together in its discussion of reshoring.[26] Physical AI is relevant where better adaptation helps a domestic operation handle changing products or deliver reliable output with the people and equipment available.

Trade and industrial policy can alter the business case from both sides. Duties can make imported goods less competitive while also raising the cost of imported equipment and components. Incentives can reduce initial investment while leaving service, skills and production costs unchanged. A buyer should therefore test equipment cost, input cost, usable demand and support as separate scenarios. An announced investment commitment should not enter the demand forecast as an accepted customer order.

In January 2026, the U.S. Department of Commerce announced a Taiwan trade-and-investment framework aimed at expanding domestic semiconductor and related industrial capacity.[27] Such sector-specific policy illustrates the environment in which automation investment is being considered. It does not establish that a particular robot, model or factory project will earn an adequate return.

Adoption also depends on local implementation capacity

Small and medium-sized manufacturers often need help selecting and integrating technology before they need a broader AI platform. Manufacturing.gov's July 2026 announcement of a funding opportunity for 14 MEP centers explicitly connects U.S. industrial competitiveness with support for smaller manufacturers adopting advanced technologies.[28] The practical route can include an external assessment, equipment trial and integrator-supported application. Public support helps develop a proposal; operations must still establish its production result and the work needed to maintain it.

The national market consequently contains different buying problems: brownfield integration, product-mix flexibility, high-rate bottlenecks and regional service capacity. A supplier that can support a local plant through its next change has a different offer from one that can deliver a model demonstration remotely. This distinction should shape procurement and the location of field-service investment.

The same machine meets three different factories

For a high-mix SME, engineering attention may be the binding constraint: every interface, recipe and recovery path consumes scarce technical time. In high-volume production, a small gain can compound across many cycles, while an interruption propagates through constrained operations. Brownfield plants add fragmented operating knowledge across controllers, instructions and experienced people. The application must identify authoritative sources, resolve disagreement and maintain those mappings when equipment or recipes change. These differences change both the first investment and the support a supplier must provide.

RoviSys industrial-AI director Bryan DeBois describes an integration stack spanning controls, networks, MES and AI in a U.S. Manufacturing Today conversation, drawing attention to the manual and operator-dependent conditions encountered in plants.[29] This practitioner perspective helps frame the integration work that a supplier must scope before promising an operating result.

The priority is to identify where scarce engineering attention and sellable bottleneck capacity justify flexibility—not to automate because a national indicator moved.

Consider two hypothetical plants. A high-mix metal SME has one controls engineer, frequent changeovers and uncertain demand. Its next step is a two-week baseline of setup, intervention and accepted output for one cell, followed by a fixture/process option and an integrator-scoped test. It should not begin with plant-wide orchestration: the pilot must reduce the work required for the next product family.

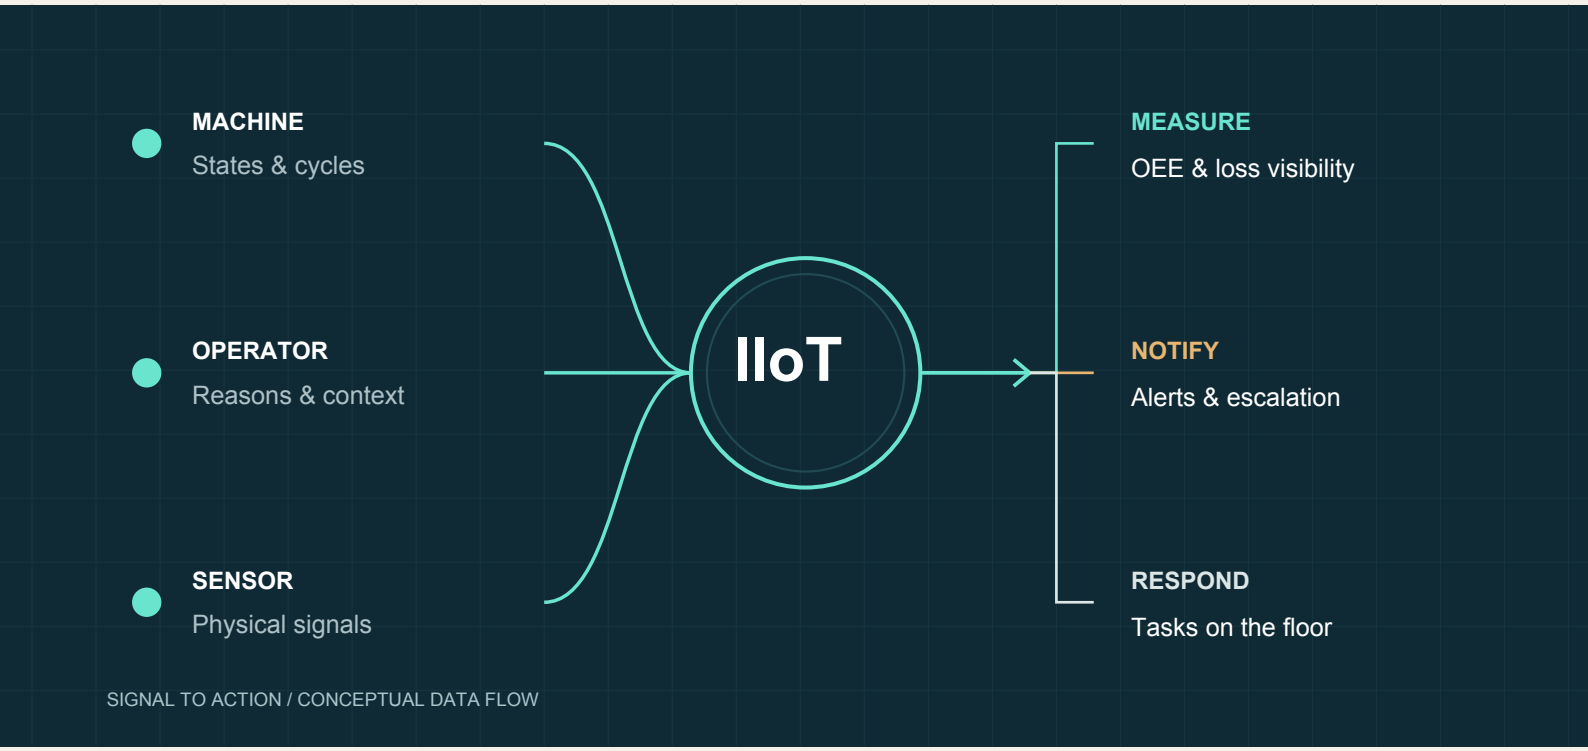
A high-volume plant has stable demand and one constrained inspection-to-assembly handoff. Its next step is different: quantify lost accepted output at that bottleneck, then compare an extension of the existing execution/control workflow with a shadow coordination test. If both produce the same decision and recovery burden, retain the incumbent path. If cross-station state resolves a measured constraint that local logic cannot, proceed to a bounded operational pilot. National utilization does not supply either plant's demand assumption.

Transferability means matching process boundary, task mix, variation, demand, controls, competence, protection, service and recovery. An unknown field is a reason to narrow the test or wait, not automatic rejection.

Industrial IoT

A machine stops. Now know why.

Connect machine signals, operator explanations and alerts to act on production losses during the shift.



SIGNAL TO ACTION / CONCEPTUAL DATA FLOW

Hardware + software. One connected stack.

Capture data from legacy and connected equipment. Add downtime reasons at the workstation. Bring status, alerts and tasks to tablet and mobile views.

Start with a defined loss and a scoped pilot on selected workstations or a line.

DBR77
INDUSTRIAL INTELLIGENCE

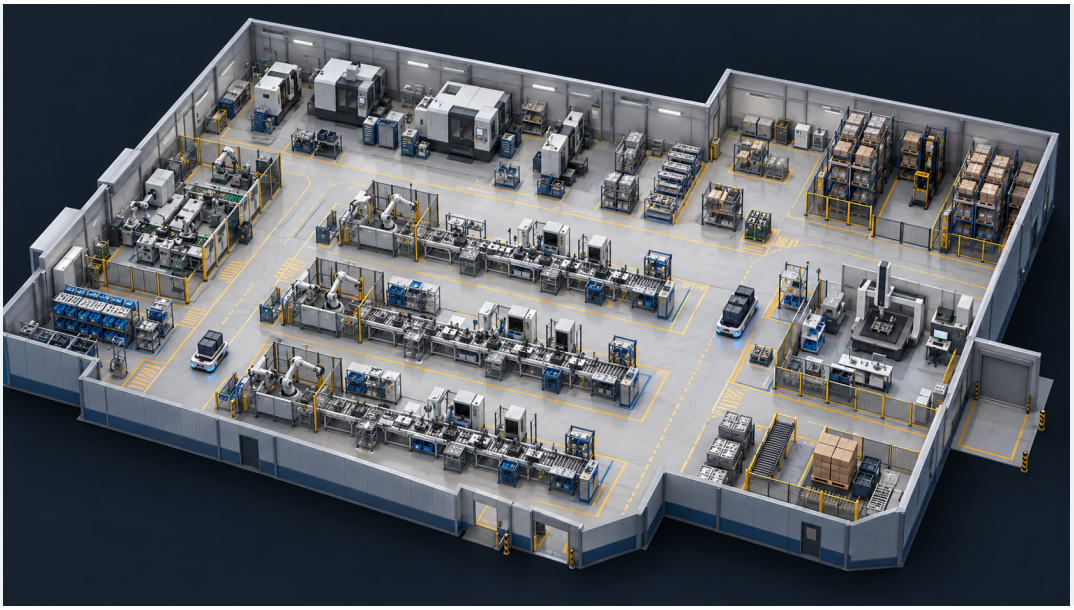
EXPLORE THE TECHNOLOGY

iot.dbr77.com

05 / FROM DECISION TO OPERATION

From decision to operation

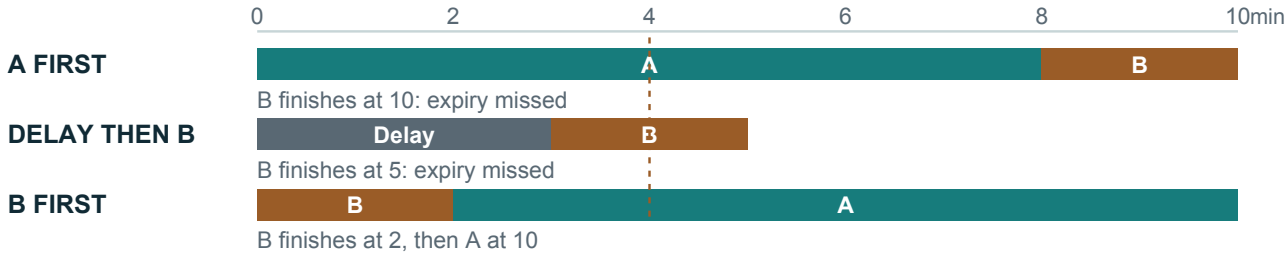
Shared operational context connects machine capability with plant performance.



AI-generated conceptual Digital Twin illustration of a fictional factory—not a customer facility or software screenshot.

SHARED INSPECTION: THREE HYPOTHETICAL SCHEDULES

Non-interruptible runs: A = 8 min; B = 2 min. B must finish by minute 4.



If B takes 5 minutes, no sequence meets expiry at 4. Do not invent feasibility.

The 3-minute energy delay is a preference. Equal information and authority; no measured uplift.

Worked example A. Three hypothetical sequences share one inspection resource. B expiry is binding; the energy delay is a preference. The incumbent rule can already choose B-first.

A recommendation, a machine movement and an execution record can disagree. Plant intelligence must reconcile authoritative state, permission to act and the physical result, rather than treating any one of those partial truths as completion.

An “operating layer” is one possible answer, not the default answer. It earns a place only if it improves a defined decision across heterogeneous equipment more effectively than extending the installed stack.

The alternative might be an incumbent automation suite, an OEM-controlled stack, an integrator-led federation, a captive system built by the operator, or a new independent layer. The report compares these routes on responsibility, not architectural fashion.

From autonomous machines to an autonomous plant

More capable machines do not automatically create a coherent production system. A robot can improve its own cycle while starving the next operation. An AMR can choose the fastest route while occupying space needed for a changeover. A maintenance model can recommend stopping an asset at the same moment a planner has committed its output to an urgent order. Each decision can be locally rational and the plant result can still deteriorate.

Consider a hypothetical conflict. Two cells need the same inspection resource. Cell A releases a large batch because its local queue is short. Cell B holds a smaller urgent order whose material expires if inspection is delayed. An energy optimizer also plans to defer the inspection station's next run. No single machine has made an obviously irrational choice; the objectives, time horizons and shared constraint were never reconciled. The plant needs a common account of state, intent and consequence—not necessarily one central controller.

Make the conflict concrete with hypothetical minutes, not operating data. Inspection is available now; A needs eight minutes, B two, and B must finish within four. Runs cannot be interrupted. The proposed energy delay is three minutes and is a cost preference, not a hard constraint. A-first finishes B at minute 10; energy-delay-then-B finishes it at 5. Both miss expiry. B-first finishes at 2, then A at 10. An existing APS or reservation rule can select B-first without learning. The new loop gains nothing by selecting the same feasible sequence at equal burden. Now suppose vision suggests a changed part condition makes B take five minutes. Give both routes that same estimate and a supervisor's staffing restriction. Neither may invent a feasible sequence or override the restriction. If the estimate is unvalidated, operations and quality own a shadow timing test; execution stays with the approved baseline or stops if that baseline is invalid. Compare accepted orders before expiry and total exception work, not prediction plausibility. Stale state or missing authority blocks both routes. This tests incremental information, not a measured Plant Operation System gain.

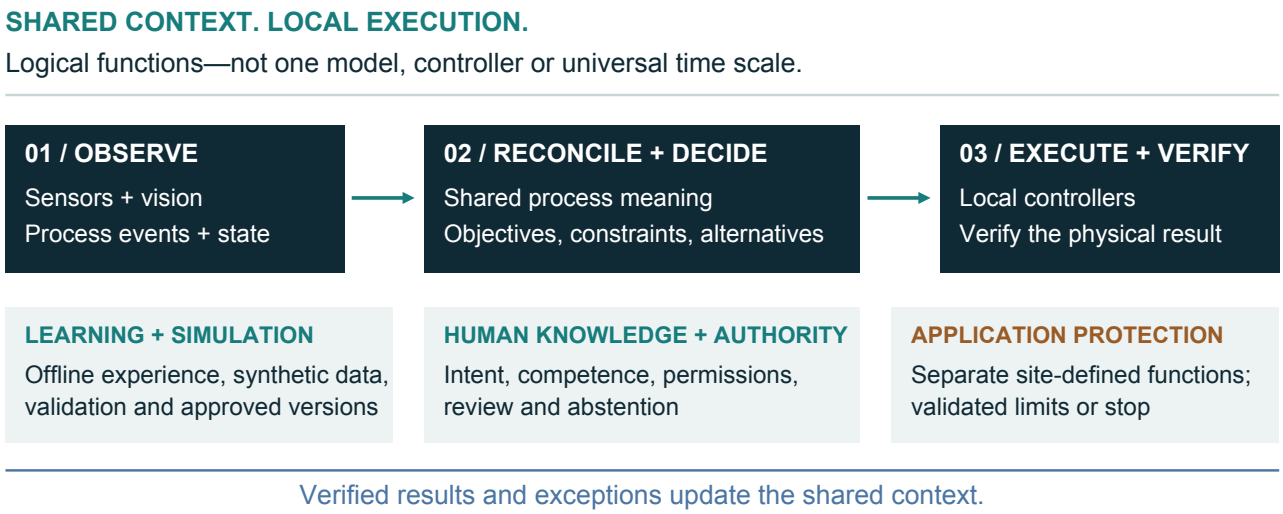


Figure 10. Shared context, local execution. Observation, decision context and local execution remain distinct. Learning, human authority and application protection constrain the process; verified results return to the shared context. Table 6 compares competing architectures.

Our hypothesis is that plant-level intelligence should connect process learning, simulation and human operational knowledge, then learn from the observed consequences of authorized decisions. Human knowledge includes practical experience, availability, workload and the reasons behind operating constraints. Capturing that knowledge is as important as collecting machine signals.

The loop is functional rather than monolithic. It observes machine, order and human-provided context; learns models within declared data and uncertainty bounds; simulates alternatives; calls an action "optimal" only relative to an explicit objective, model, constraints and conditions; delegates only authorized decisions; and records physical readback before expanding autonomy. It does not claim a globally best answer while hiding uncertainty or conflicting interests. Synthetic data may support a selected training or test task, but simulation credibility and transfer to real operation require separate gates. Feedback does not imply uncontrolled online retraining.

"One context" does not mean one model or computer. Training can occur offline, inference near the machine, and plant planning on a longer horizon. Basic control and functionally parallel protection remain local to their requirements. Marketplace activity can help procure and compare capabilities, but it is an ecosystem function, not a runtime controller.

DBR77 is building toward this vision of an autonomous Plant Operation System. It brings together perception and industrial data, process learning, Digital Twin simulation, synthetic data and operational knowledge supplied by people. The intended benefit is better coordination across heterogeneous assets and changes in production. The development test is improvement in order performance, quality-approved output or exception cost after integration and support are included.

Table 6. Competing routes to plant coordination

Architecture	Natural advantage	Principal risk	Proof that should decide
Incumbent suite extension	Installed base, semantics, service and controls access	Lock-in or incomplete cross-vendor view	Same-decision improvement and extension cost
OEM-controlled stack	Tight endpoint integration and warranty boundary	Portability and cross-plant neutrality	Accepted task result, interface rights and exit path
Integrator-led federation	Brownfield adaptation and vendor choice	Bespoke knowledge and recurring project cost	Falling effort on second and third deployment
Captive operator platform	Proprietary process data and operating discipline	Weak transfer outside owner network	Comparable result in a materially different site
Independent operating layer	Cross-asset context and potential reusable orchestration	New dependency, data access and responsibility gap	Incremental outcome over installed stack, recovery and lifecycle economics

Authors' competitive architecture framework. One supplier can provide several boxes; functional boundaries do not require separate vendors.

Five credible architectures

The incumbent suite begins with installed controllers, engineering tools, MES/MOM and service relationships. Its advantage is access: plant semantics, lifecycle support and change procedures may already exist. Its weakness is that cross-vendor context and new decision loops can remain fragmented.

An OEM-controlled stack binds intelligence closely to the robot or machine. It can simplify responsibility for motion, warranty and updates. It can also constrain portability and make plant-wide coordi-

nation subordinate to an endpoint vendor's interfaces.

An integrator-led federation joins existing systems around a use case. It respects brownfield reality and can combine best-of-breed components. The danger is that crucial knowledge remains in project work, scripts and individuals rather than a repeatable product.

A captive operator platform is built around one company's process and scale. It can exploit proprietary data and operating discipline. Its evidence may be powerful internally but difficult to transfer to an SME or a plant with different equipment and support capacity.

An independent operating layer seeks to maintain context and coordinate decisions across endpoints. Its promise is neutrality and reuse. Its burden is severe: it must obtain trustworthy data, integrate without destabilizing local control, define authority, survive platform loss and show that the incremental decision is worth another lifecycle dependency.

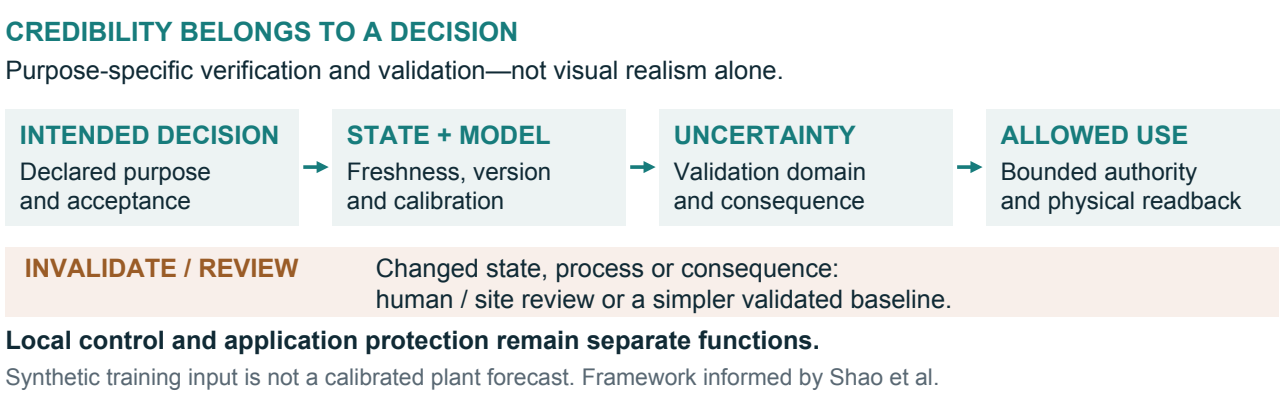


Figure 11. Digital Twin credibility. A model's credibility belongs to its intended use. Data quality, model validity, uncertainty and change determine the decisions it can support.

What a twin is credible for

A digital twin can support several different jobs: engineering visualization, virtual commissioning, state estimation, scenario analysis, maintenance or operational recommendation. A model credible for one is not automatically credible for another. A three-dimensional display may help people inspect layout. It does not by itself predict queue behavior, contact physics or the consequence of a changed recipe.

Shao, Hightower and Schindel's manufacturing-twin framework treats credibility as an assessment tied to intended use and maintained through the model's lifecycle. It includes data, implementation, operational validity and uncertainty in the resulting decision.[11]

That boundary turns "live" into a testable claim. Which signals are connected? How fresh are they? Does the model converge with authoritative physical readback after change? What uncertainty reaches the decision? Which event invalidates the model? A twin can be synchronized and still be wrong for the use that matters.

For a scheduling decision, a simpler queue model may beat a detailed twin. For a collision-sensitive manipulation study, task-relevant geometry and dynamics may matter more than visual fidelity. For maintenance, the model must distinguish a useful warning from an alert burden that operators learn to ignore. The correct baseline is the least complex representation that changes the decision reliably.

A worked trace without an omnipotent model

Return to B's inspection reservation. Cameras, PLC tags and the supervisor's declared restriction supply different inputs; order expiry and equipment mode constrain the same choice. A service reconciles their timestamps and provenance before recommending B-first, hold or escalation. The recommendation has no physical authority by itself.

The operational route is: observation, state estimate, context reconciliation, recommendation, authority decision, bounded command, basic control, actuation, physical readback and acceptance or recovery. Each transition can fail differently. A stale tag can corrupt context. A plausible recommendation can violate the current mode. A valid command can execute while the expected part is absent. A command acknowledgement can arrive even though the process result is unacceptable.

The protection route is functionally parallel. It observes the safety-relevant signals established by the application and can limit or stop hazardous motion regardless of the AI recommendation. It does not blindly execute model instructions, and the AI cannot certify, suppress or reset it. “Independent” describes the degree of functional independence and common-cause control required by the validated application; it does not require that every function come from a different vendor or physical box.

The human role must be equally specific. Approval is meaningful only if the person has the information, competence, authority and time to act. A supervisor can reasonably review a reversible scheduling choice over minutes. The same interface is not a safety case for fast hazardous motion. When the model abstains, the system must identify who receives the case and what the plant does while waiting.

Readback asks whether B passed inspection before expiry and A’s subsequent work was accepted, not whether software sent a reservation. Observed duration and the operator’s explanation of the changed condition become evidence for a reviewed model update, not automatic retraining. If exceptions require repeated manual verification, better planning may still fail to reduce total work.

When the incumbent is enough

A new layer is unnecessary when the installed controller, MES, WMS or APS already has the required state, authority and readback at acceptable lifecycle cost. It is also unnecessary when a local deterministic rule resolves the decision, or when the variability can be removed upstream. Architectural elegance is not a business benefit.

The case for a new layer strengthens when several heterogeneous systems must share context, the same decision pattern recurs, interfaces are maintainable, and the improvement survives the cost of integration and exception handling. It strengthens further when local operation remains valid after platform loss, versions are traceable and the customer retains an exit path.

It weakens when the platform needs uncontrolled endpoint access, duplicates records, creates common-mode failure, or hides human service behind “autonomy.” Vendor count is not the principle; bounded authority and accountable operation are.

In a September 2026 Siemens/ARC conversation, Rainer Brehm described the installed-base instinct as “never change a running system,” then explained the role of testing and virtual validation in making updates credible.[14] This captures a practical adoption problem: the value of a change must exceed both its implementation cost and the risk of disrupting dependable work.

For DBR77, the priority is to demonstrate that shared context can be reused across recurring decisions. A credible operating product combines that reuse with maintainable interfaces, a clear division of responsibility and a measurable plant result.

The deployment contract

Before a pilot, management should require a compact operating contract: the decision, affected equipment, maximum authority, authoritative inputs, abstention behavior, independent protection, readback, reset and restart roles, degraded baseline, change control, economic owner and the stop/continue/expand rule. Exact thresholds belong to the site and remain unknown until approved.

This contract reveals whether the product is a recommendation tool, a bounded execution system or a claim with no owner. If the plant cannot authorize or verify the action, the missing integration is the product.

ADVERTISEMENT / DBR77

Digital Twin

Test the change. Not your factory.

Compare production scenarios before committing capital to changes on the factory floor.



A better decision starts with a better model.

LAYOUT

Explore alternative factory configurations.

FLOW

Examine material and resource movement.

SCENARIOS

Compare assumptions and operating choices.

AI-generated fictional factory concept. Not a customer site or software screenshot.

DBR77
INDUSTRIAL INTELLIGENCE

EXPLORE THE TECHNOLOGY

dt.dbr77.com/en

06 / THE ECONOMICS OF EXCEPTIONS

The economics of exceptions

When flexibility pays — and when a simpler process should win

The normal cycle is the easiest part of an automation proposal to sell. A part moves, a task completes, a dashboard records success. The economic question begins with everything outside that sequence: the parts the system declines, the assistance it needs, the time required to restore work, and the capacity the business can actually sell. Physical AI may reduce the cost of handling variation. It can also move that cost from a fixture into integration, data preparation, supervision and support. The investment case must identify which of those things has happened.

Comparing alternatives makes uncertainty manageable. A manufacturer can rationally choose a new AI capability, a conventional cell, an extension of its existing system or a process change. A useful proposal identifies the conditions under which each would be preferred and the missing measurement worth obtaining before commitment.

What the public cases can tell us

NIST MEP's July 2024 account of AMG Industries in Ohio describes a loaned cobot tested on a repetitive exhaust-tip task. It reports output increasing from 200 to 276 parts per hour, followed by purchase of a robot. This is a positive program/customer account, not an independent lifecycle-cost audit. It does not establish the contribution of AI inference, nor provide a complete intervention and recovery denominator.[30]

The MEP account of Michigan job shop JCGibbons, published in October 2022 and updated in June 2023, likewise describes assessment, a loaned cobot and integration support before the company's purchase. Its relevance is the path to an informed acquisition, not proof of today's generalist AI economics. The public story does not supply a complete comparable cost and operating record.[31]

These cases make a useful counterpoint to a technology-first pitch. The route into automation can include a bounded task, borrowed equipment and external expertise. But support is an input, not magic. A real comparison must establish who paid for that work, what remains after the trial and what another factory would have to reproduce. A positive account can justify investigation without filling the missing fields of the buyer's own business case. Neither story supplies a category-wide return on Physical AI.

Brian Weinstein of Connecticut manufacturer Chapco describes automated sheet loading, unloading and sorting alongside experimentation with AI in ManufactureCT's publisher transcript. Skilled setup and programming remain part of that operating model.[32] His account illustrates how automation can release human attention for work that still requires experience.

Removing repetitive handling and releasing skilled time can be valuable without eliminating a job. For the business case, establish the actual disposition of that time: more saleable work, avoided overtime, a vacancy not filled, or simply unused capacity. Do not claim the full wage bill as cash saved when payroll does not change. Benefits to ergonomics or resilience deserve explicit evaluation, but a convenient monetary number should not be invented to make an incomplete return look positive.

One task, four credible choices

Return to the hypothetical task of presenting a part despite variable position. We compare the current process, a fixture or local process improvement, an extension of installed workflow/control with an integrator, and a new local AI policy package. This last option is **not** a plant-wide operating system. Faster manipulation cannot establish the value of a separate coordination layer.

The calculation below is deliberately synthetic. H is 1,000 normalized time units per planning period. An attempt is one logged execution opportunity within the agreed task boundary. q is the fraction of all attempts delivering an accepted result, fixed at 0.95; e is the fraction of that same population requiring recovery. Failed, aborted and refused attempts remain in the denominator. c covers time per attempt excluding separately counted recovery; r is additional time per recovery event. A simple planning estimate of accepted output is:

$$G = Hq / (c + er).$$

This approximation treats average cycle and recovery times as stable, with no overlapping recovery. It is not a finite-window stochastic simulation. Correlated failures, queues, changeovers and rare long outages require a richer model. Recovery time consumes capacity here; its separate labor or service expense belongs in operating cost, without counting the same financial loss twice.

Table 7. Assumed configurations and incremental costs

Option	c	e	r	G per period	K	C per period	X
O0: current process	1.00	0	0	950.00	0	0	0
O1: fixture/local improvement	0.90	0.01	1	1043.96	40	15	4
O2: incumbent + integrator	0.85	0.02	2	1067.42	70	20	6
O3: new local AI package	0.80	0.01	3	1144.58	100	40	10

Authors' synthetic assumptions. K: initial cost; C: recurring cost; X: exit cost. Costs are normalized value units, not dollars. O0 is an incremental baseline, not free production. Different c/e/r are hypotheses about configurations, not measured effects of software architecture.

K includes setup, integration, ramp, validation and training; C includes incremental supervision, support, compute, maintenance and other recurring expense; X covers exit. They must be itemized without overlap in a real proposal. Contribution per additional accepted unit, m, is normalized to one after unit-variable costs; those same costs must not reappear in C. No additional saving from the same labor hours is added to this example.

Turning capacity into economic value

The demand limit is the number of additional units the business can use or sell beyond its current output. The benefit is the contribution from the smaller of that demand and the capacity gain. Recurring support and supervision costs are then deducted, followed by initial and exit costs over the three-period illustration. A reduction below baseline is counted as lost contribution. The equations and input definitions appear in the technical note after the main chapters.

There is no discounting, tax, financing or residual value here: V is not NPV. The purpose is to trace one common choice from technical assumptions through usable output to cash. A real investment model needs period-specific ramp, demand, working capital and the company's finance assumptions. Normalizing units avoids inventing prices; it does not make unknown commercial inputs known.

Table 8. The same choices under different demand limits

Additional demand D	V: O0	V: O1	V: O2	V: O3	Conditional economic winner
0	0.00	−89.00	−136.00	−230.00	O0
40	0.00	31.00	−16.00	−110.00	O1
120	0.00	192.87	216.25	130.00	O2
220	0.00	192.87	216.25	353.73	O3, subject to its assumptions

Authors' calculation from Table 7, before rounding intermediate output. The same task, quality, demand cap and horizon apply to every option. This is not a recommendation to purchase or operate equipment.

Read O3 at D=120 from left to right. Its estimated output of 1,144.58 exceeds baseline by 194.58, but only 120 additional units can be sold. With m=1, benefit B is therefore 120, not 194.58. Deduct recurring cost of 40 to obtain F=80. Three periods yield 240; subtract initial cost 100 and exit cost 10 to obtain V=130. O2's V of 216.25 is higher despite its lower technical capacity. O3's additional capacity is only partly used under this demand assumption.

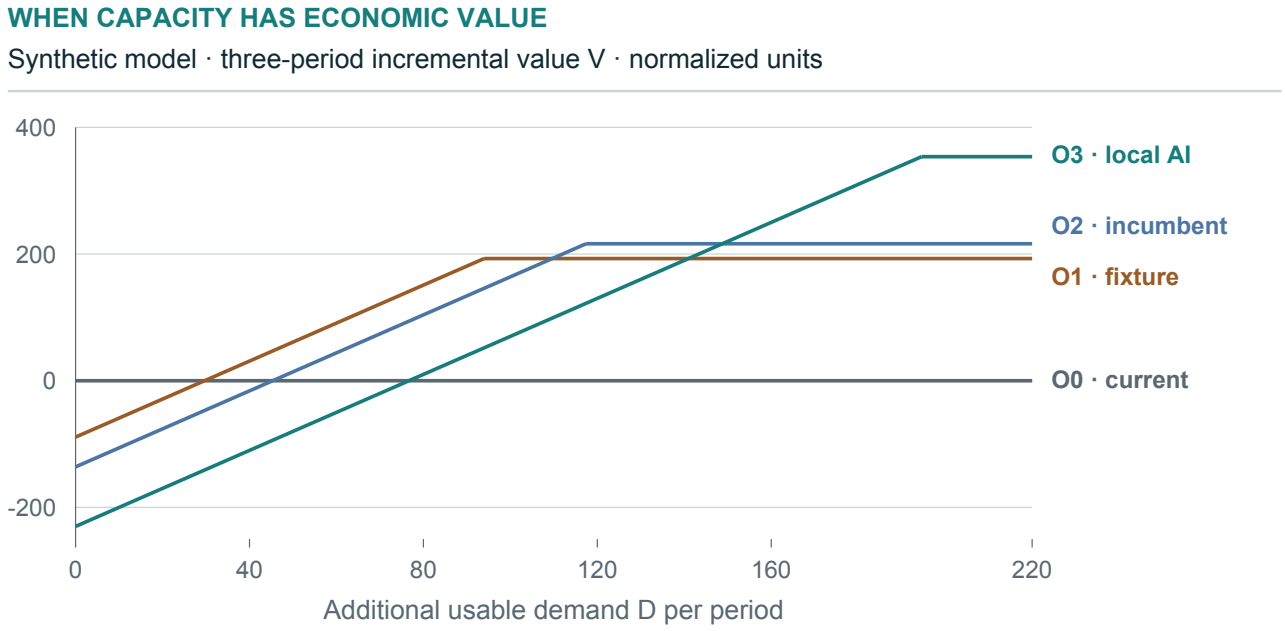


Figure 12. Demand changes the preferred option. Synthetic model with fixed performance and cost assumptions. Higher technical capacity creates value only where additional accepted output is usable. The four options are compared under the same demand limit and three-period horizon.

Every option can win. With no additional demand, the extra capacity has no modeled benefit. At modest demand, the least expensive improvement is enough. At D=120, extending the incumbent

system wins. Only with sufficient usable opportunity does the more capable local AI configuration justify its higher assumed cost. A smaller manufacturer is not inherently the low-demand case; a large factory is not inherently the high-demand case. For high-mix work, repeat the comparison across the relevant mix and changeovers rather than trusting one average part.

The illustration deliberately leaves some possible sources of flexibility value unpriced. A new capability might reduce changeover effort, protect delivery dates or make a short run feasible; those are separate hypotheses to measure, not free additions to B. Compare the same orders and service commitment, including conventional alternatives that can also handle the mix. Where the commercial benefit is the ability to respond to uncertain future work, disclose the demand scenarios and carrying cost of keeping that option available. Do not call an attractive possibility contracted demand.

Where the attractive answer reverses

At D=220, increase O3's recovery-event share from 0.01 to 0.05, leaving its other inputs unchanged. Output falls to 1,000 and its three-period net becomes −80. O2 now wins. Alternatively, retain O3's original performance but increase recurring cost from 40 to 100: its net falls to 173.73, again below O2. The same example captures two different problems—lost useful capacity and expensive support—without inventing another source of benefit.

With other inputs fixed, O3 overtakes O2 at roughly 150 additional demand units. At demand of 220, O3 remains ahead only while recovery events affect fewer than about 2.2% of attempts. These are rounded economic thresholds within this synthetic model. They are separate from the application's safety requirements. The unrounded calculations are retained in the technical note.

If e is unknown between 0.01 and 0.05, selecting its favorable endpoint would manufacture the recommendation. The useful answer is to defer commitment to O3 and obtain a bounded record of attempts, accepted results, assistance and recovery. Operations and quality establish performance; finance confirms demand and the cost perimeter; controls, EHS and IT/OT determine whether and how a physical trial is authorized. Without probabilities, the example does not calculate expected value of information or prescribe a trial budget.

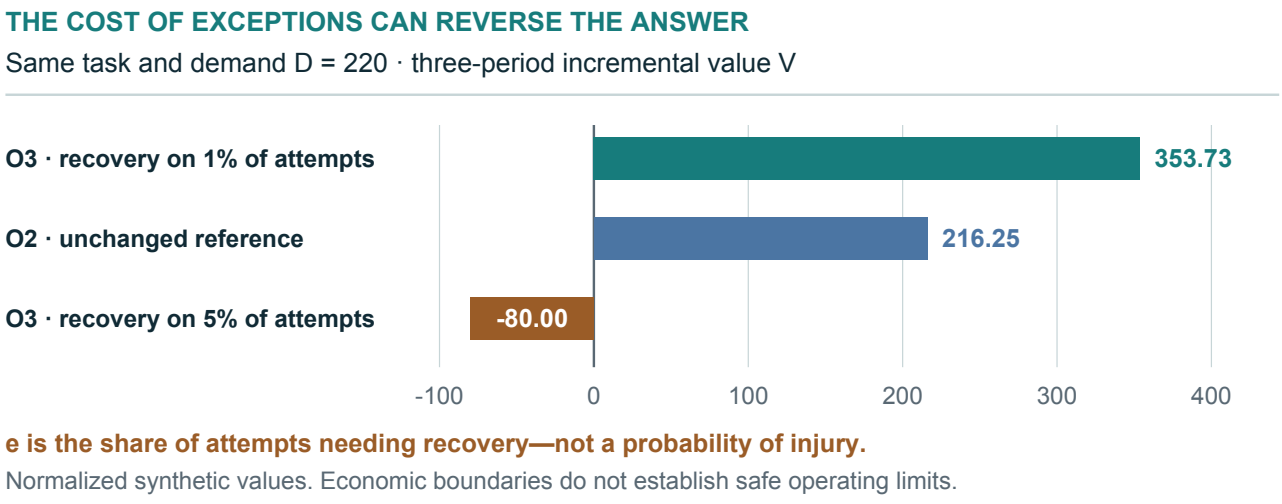


Figure 13. Recovery changes the economic winner. At D=220, changing only O3's recovery-event share from 0.01 to 0.05 reduces modeled value from 353.73 to −80.00. O2 remains at 216.25. Synthetic normalized values—not measured plant returns or safety thresholds.

The supplier has a different ledger

On Alpha Partners' July 2026 podcast, Schematic Ventures founder Julian Counihan discusses operating relationships and physical-asset financing. Describing the shift from software uncertainty into physical industries, he says, "you can't vibe-code a factory or an actuator or a motor."[33] The implication for a supplier is to account for the equipment, engineering and service needed to fulfill its promise.

A buyer's benefit and a supplier's economics are not interchangeable. In a separate synthetic supplier index, revenue of 100 faces equipment/depreciation 20, field engineering 15, service 12, remote work 10, compute 5, warranty/spares 5, financing 8, success support 5 and lifecycle validation 5. The remaining contribution is 15 before R&D, overhead and tax—not net margin. An additional 10 of remote work leaves 5; an additional 20 of engineering turns contribution negative.

Cash can fail before that contribution is realized. If initial supplier funding is 60 and only 20 is recovered before start, the funding gap is 40. That gap is not another expense to deduct from the same contribution. Investors should ask how it changes across deployments, alongside actual collections, support effort and renewals. A recurring contract is not by itself evidence of low-touch software economics. Nor does an upfront sale establish that the manufacturer owns the knowledge required to maintain the result.

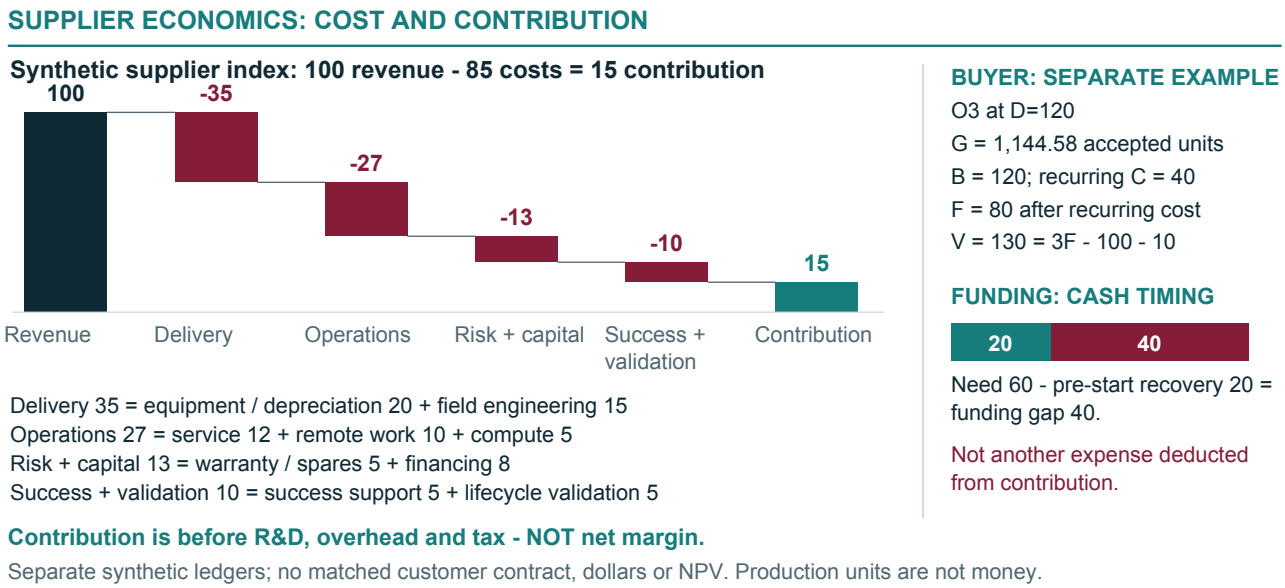


Figure 14. Buyer value and supplier contribution. The synthetic buyer example connects usable output with cost. The supplier example separately traces revenue, contribution and initial funding.

Buy the next justified step

The comparison should end with a decision and its missing evidence. Choose the simpler alternative when it meets the required result at lower full cost. Investigate a new capability when the remaining variation has demonstrable value and its economics depend on a measurable uncertainty. Stop physical execution if authority, protection or recovery is undefined. Scale only after the next operating context and its support burden have been tested against stated conditions.

Coordination across several workstations creates another source of value: better routing, fewer conflicting decisions and more dependable order completion. The corresponding business case should measure those plant effects separately from the manipulation gains illustrated here. This gives DBR77's proposed operating layer a clear performance agenda and allows customers to compare its contribution with an extension of their existing systems.

Marketplace

Buy a solution. Not a sales pitch.

Turn one automation challenge into comparable supplier proposals - with scope, assumptions and exclusions visible.

ONE CLEAR AUTOMATION BRIEF

The task. The process. The constraints.

01

SCOPE

What is included?
Does the solution
fit the task?

02

ASSUMPTIONS

What must be true?
Which inputs and
dependencies matter?

03

EXCLUSIONS

What is not covered?
What could change
cost or delivery?

A common basis for comparison. Not three unrelated pitches.

Define the challenge, engage technology suppliers and integrators, and compare proposals against the same operational need.

Illustrative procurement framework; not live supplier offers.

DBR77
INDUSTRIAL INTELLIGENCE

EXPLORE THE TECHNOLOGY

marketplace.dbr77.com

07 / WHO CAPTURES VALUE?

Who captures value?

The winning layer is the one that earns authority, carries exceptions and retains economics—not the one with the most expansive category label.

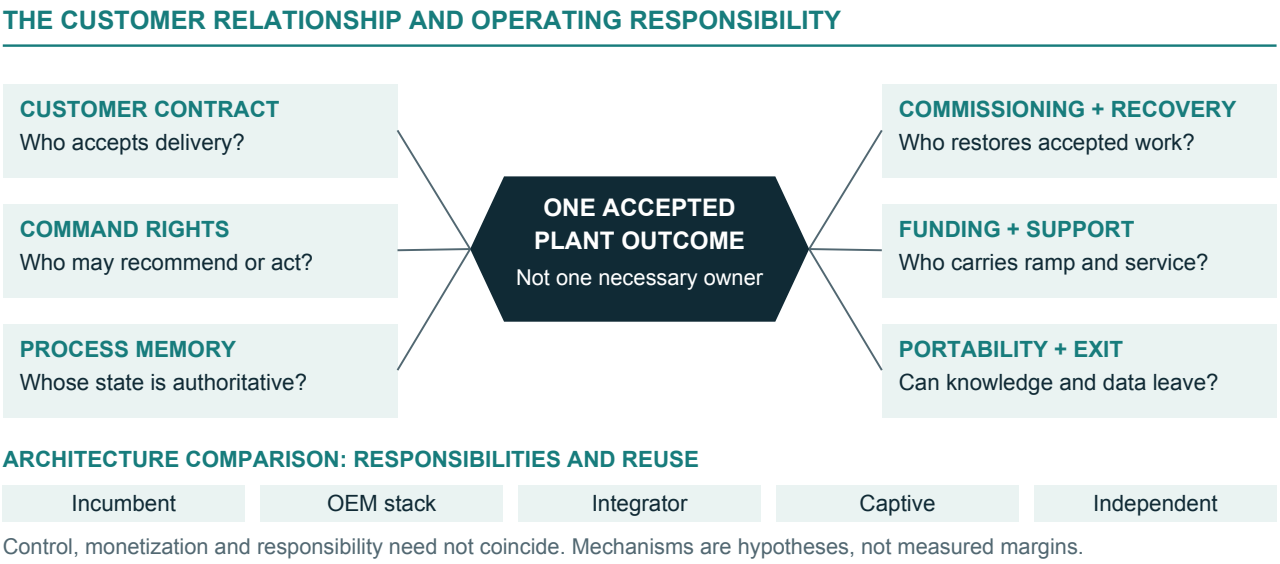


Figure 15. The operating relationship. Contract ownership, process knowledge, decision rights and service responsibility shape how industrial value is captured.

Physical AI can create value at a task while sending revenue, data advantage and the customer relationship elsewhere. Model provider, OEM, automation incumbent, integrator, service provider and manufacturer may control different parts of the result. The strategic question is who holds the right to act, the obligation to recover and the route to the next deployment.

Follow five control points

Five control points matter: customer contract; authority to recommend, schedule or command; process meaning across orders, recipes, machine state and quality; exception handling; and portability without rebuilding operating knowledge.

These control points can belong to one company or several. Functional separation need not fragment supply, and one prime contractor does not eliminate internal boundaries. Responsibility must remain legible when normal operation stops.

Symbotic's FY2025 filing describes integrated hardware-and-software systems that it initially operates before transferring daily operation to the customer; Walmart accounted for about 85% of FY2025 revenue, while Walmart and GreenBox represented the vast majority of backlog.[34] A filed presentation separately reported 48 operational systems, 50 systems in deployment and 11 customers at fiscal year-end.[35] These are issuer-reported commercial categories, not an independent census of accepted performance or customer return, but they show why contract concentration, deployment duration and the distinction between operating and deploying systems belong in platform diligence.

Teradyne reported more than 110,000 cobots and 11,000 AMRs sold globally, yet its Robotics revenue fell 15.5% in 2025 and the segment recorded restructuring and an inventory provision associated with revised demand forecasts.[36] Installed or sold units can coexist with cyclicality, channel depen-

dence and capacity adjustment. Scale of installed hardware is not itself proof of current demand, active utilization or recurring software economics.

The portfolios of Rockwell Automation and Siemens illustrate an incumbent countermodel: devices, control/software and lifecycle services can be monetized together through an installed base.[37, 38] Their reported segment revenue is much broader than Physical AI, so it cannot size the category. It does show that a new operating layer competes against distribution, service, controls and switching costs—not against a blank architecture.

"Platform" is not one business model

A platform may be a reusable product, engineering organization, channel strategy or label over project work. Each can create value with a different cash profile.

In a publisher transcript, Automated Industrial Robotics CEO Darragh de Stonndún describes growth through broader engineering capability and responsibility for more of the customer's automation problem.[39] This is a substantive industrial platform strategy. Its repeatability can be tested through the engineering hours, support effort and contribution of successive comparable deployments.

For a technology company, build the task, data or interface asset that becomes more reusable with each accepted deployment. Buy commodities when ownership adds no differentiating knowledge. Partner where commissioning, controls, protection, service or channel access would take too long to recreate. Do not call repeated service software when its labor does not decline; do not dismiss service when it makes the outcome dependable.

Former Amazon operations executive Jeff Wilke makes the ownership problem plain: "If you've outsourced all of your technical thinking, it's going to be way harder for you to make this transition."[40] A manufacturer can use specialist suppliers while retaining the competence to challenge decisions and improve its operating system.

For an investor, bridge contracted interest to operated cohorts. Separate hardware, licenses, engineering, commissioning, remote work, service, warranty, financing and validation. Ask which costs precede acceptance, which recur, who funds equipment, whether support per unit falls, and whether renewal follows accepted work. Backlog can absorb working capital; recurring contracts can conceal intensive support. Both change valuation logic.

The business behind plant-level intelligence

DBR77's strategic wager is that reusable operating knowledge can connect local intelligence with coherent plant decisions. Value could come from recognizing cross-asset constraints sooner, testing responses in simulation and transferring the resulting decision pattern to another operation. These are measurable product objectives, rather than benefits inferred from the breadth of a portfolio.

Chapter 5 defines that loop and its limits. Here the issue is who can turn it into a business. The following strategic reading links named participants to the control point and the burden that could consume its return; it does not claim measured margins for these approaches.

Table 9. Control points and monetization

Participant / route	Potential control point	Monetization mechanism	Burden that can consume the return
Symbotic's integrated-system route	System delivery and customer operating relationship	Integrated deployment and lifecycle offering	Concentrated contracts, ramp and deployment cash
Teradyne's cobot/AMR portfolio	Endpoint and route to application	Equipment and associated ecosystem	Demand cycles, channel and integration dependence
Rockwell / Siemens incumbent route	Installed engineering, controls and lifecycle relationship	Bundled hardware, software and service	Legacy interfaces, cross-vendor integration and change cost
AIR / integrator-led route	Responsibility for the customer's automation problem	Engineering, integration and service	Bespoke effort that fails to become reusable
Independent layer, including DBR77's proposed direction	Portable cross-asset context and decision outcome	Product or managed-outcome proposition to be proved	Permission, integration, exception service and another dependency

The incumbent advantage may be less a better model than cheaper access to trusted state, command permission and service. An independent layer wins only where it resolves a valuable cross-vendor decision that the installed route cannot resolve at comparable lifecycle burden. A marketplace may improve sourcing without owning that runtime control point. A twin may improve a decision without becoming the system of record. Portfolio breadth is therefore not the moat; reusable operating knowledge, permitted access and a defensible outcome relationship are candidates for one.

The route should follow the work that remains. Extend the installed APS where its rules and interfaces suffice. Use an integrator where trustworthy inputs need to be connected. Evaluate a reusable operating layer where the same cross-asset decision recurs across different equipment. In each route, contracts should identify the prime contact for exceptions and preserve exportable event history, mappings and configuration. Manufacturers retain responsibility for production objectives and quality acceptance; suppliers earn recurring value by reducing the effort required to deliver and maintain the result.

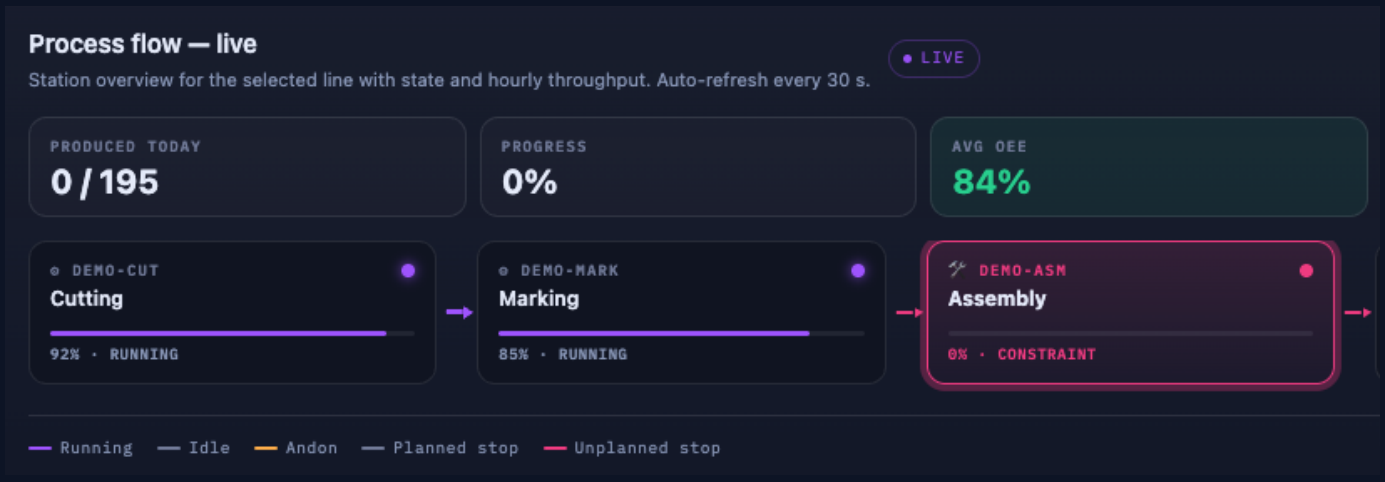
ADVERTISEMENT / DBR77

IRIS / Plant Operation System

One plant.
One operating picture.

Connect production, quality, maintenance and tasks -
so teams can act on the same operational information.

IRIS / PROCESS FLOW



SEE

The exception.

Shared production context and process status.

ASSIGN

The next action.

Tasks and responsibilities across functions.

FOLLOW THROUGH

The resolution.

Workflows connecting information to execution.

AI recommendations support decisions. People retain approval authority.

Actual IRIS demo detail. Displayed metrics are example data, not customer results.

DBR77
INDUSTRIAL INTELLIGENCE

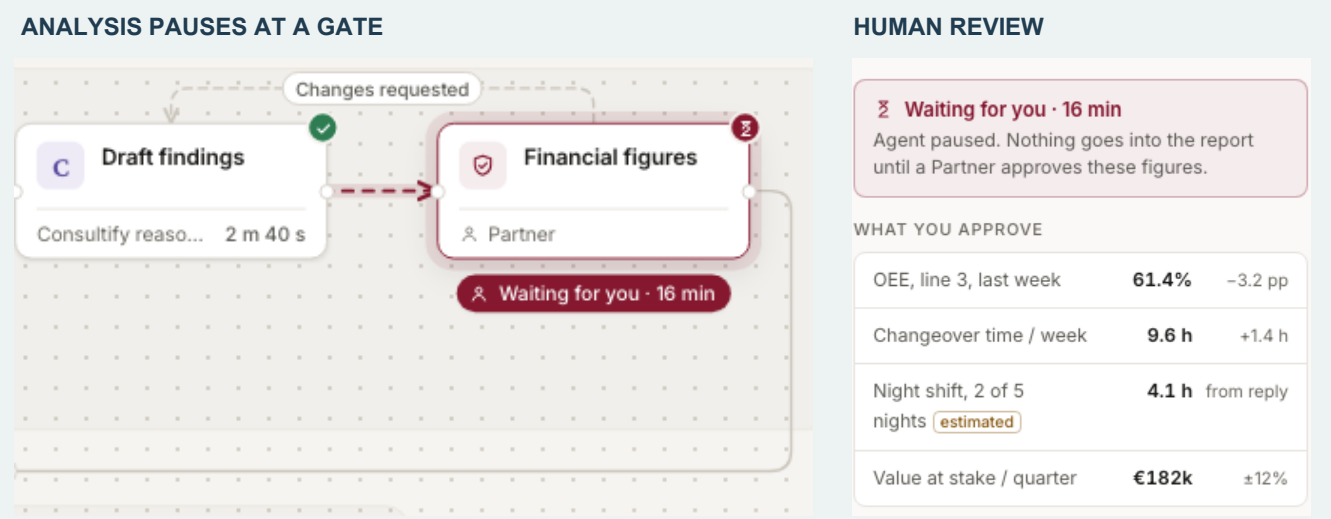
EXPLORE THE TECHNOLOGY

iris.dbr77.com

08 / WORK AND ORGANIZATIONAL CHANGE

Work and organizational change

Productivity belongs to a unit of analysis. Responsibility belongs to people even when a system recommends the action.



Consultify illustrative agent workflow; displayed values are examples. Analysis pauses at an approval gate. Generating a report does not confer authority to control the plant.

Table 10. Task, competence and voice

Work element	Possible change	Evidence to retain	Organizational decision
Setup and changeover	Less manual programming or new data preparation	Time, skill and errors before/after	Who owns recipes, prompts, policies and validation?
Normal execution	Automated movement or recommendation	Accepted output, intervention and exposure	Which task is removed, retained or shared?
Exception and recovery	New diagnostic or teleoperation work	All assists, stops, restart time and handoffs	Who is trained, available and authorized?
Quality and oversight	More traceability or more monitoring	Rejections, overrides, false alarms and appeals	Who can challenge the system and correct records?
Maintenance and improvement	New software/data dependencies	Support hours, version changes and retained know-how	What competence stays at the plant?

A plant can become more productive while a particular occupation, shift or community bears the adjustment. Automating repetition can also increase setup, maintenance and exception work. Ask which tasks change, whose output improves and who can contest a poor result.

Productivity, employment and distribution

Acemoglu and Restrepo studied U.S. commuting-zone exposure to industrial robots, principally over 1990–2007, and estimated adverse local employment and wage effects.[41] The paper distinguishes local from implied aggregate effects and cautions against simple future extrapolation.[41] It shows why geography and distribution matter, not how to staff a 2026 plant.

Graetz and Michaels used data for 17 countries over 1993–2007 and reported productivity gains without a statistically significant reduction in aggregate hours, alongside deterioration in lower-skilled workers' relative position.[42] The result shows that productivity, hours and distribution can move differently; it does not guarantee a harmless deployment.

The studies use different units and periods. A plant should track accepted output, hours by task, training, redeployment and work shifted to contractors—not only headcount.

Redesign tasks before announcing jobs

Physical AI work includes preparation, execution, supervision, exception diagnosis, recovery, maintenance, quality acceptance and improvement. A deployment may augment one task, displace another and create a third. The outcome depends on role design, not only model capability.

Consider an explicitly illustrative cell where automated loading removes two hours of repetitive handling per shift but moves setup, alarm review and restart work to one experienced operator. Production gains time; that operator gains responsibility and monitoring load, while newer operators lose practice diagnosing the process. The conflict appears when management expects the expert to cover two cells without protected recovery time or authority to reject a dubious restart.

The organizational decision is concrete: before scale, assign trained backup coverage, reserve paid training and exception-review time, log interventions by shift, and give the designated operator stop-and-escalate authority without a production-target penalty. Scale only if accepted output rises without unresolved recovery queues or dependence on one person. Otherwise retain one cell and redesign the role. This is an analytical example, not a customer story or measured outcome.

An April 2026 survey prepared by David Binder Research for the AFL-CIO reported strong preferences for human final employment decisions, training and worker input into workplace-technology design among an analytic group of 1,164 registered voters who were employed or looking for work, within a 1,588-person design with a union-member oversample.[43] These are stated preferences from a sponsored survey—not manufacturing adoption, safety performance or evidence that every machine command needs manual approval.

In her April 2026 National Press Club speech, AFL-CIO President Liz Shuler calls for worker participation, training, privacy and retained human responsibility, with collective bargaining and policy as implementation mechanisms.[44] That perspective adds an important operating question: who has the knowledge and practical authority to shape how new technology changes work?

Participation is operational. Operators and maintenance staff often know which variation is normal and which recovery creates downstream trouble. Their knowledge should enter task definition, exceptions, acceptance and review. Consultation does not replace engineering or protective functions.

Training must follow responsibility. Supervisors need time, information and authority; recovery roles need diagnostic access and practice. Scale requires evidence about new work, monitoring, retained competence and fallback. “Reskilling” without a named role, curriculum, time and employment path is not an operating plan.

Physical AI may expand output and remove unpleasant work. It may also centralize monitoring or shift exceptions to contractors. The evidence supports neither a ten-year unemployment number nor an absence-of-harm claim; it supports measuring distribution alongside output.

ADVERTISEMENT / DBR77

Consultify

AI implementing AI.

Delegate the transformation work. Keep the decisions that matter.

AI agents turn company context into analyses, business cases, roadmaps and working initiatives - not just another slide deck.

01 Diagnose & prioritize

Read company sources. Structure readiness and identify transformation opportunities.

DELIVERS **Assessment + opportunity backlog**

02 Build the case

Develop business cases, prioritize the roadmap and prepare management materials.

DELIVERS **Business case + roadmap**

03 Run the transformation

Organize initiatives, owners and milestones. Produce working documents and track progress.

DELIVERS **Initiatives + progress reporting**

People approve key decisions. Technical integrations and deployment scope are agreed separately.

DBR77
INDUSTRIAL INTELLIGENCE

EXPLORE THE TECHNOLOGY

consultify.ai

09 / RISKS AND RECOVERY

Risks and recovery

A physical system is judged not only by the action it chooses, but by the path from error to harm—and the path back to legitimate operation.



AI-generated conceptual illustration.

Physical AI extends the consequences of software decisions into equipment and production. An incorrect inference can move a tool, damage material, disrupt flow or expose a person. Evaluation therefore follows the entire path from observation and uncertainty through the authorized action to its physical result and recovery.

The positive case for AI is important. Better perception can reduce awkward manual handling. Adaptive systems can take on variable or ergonomically demanding work. An operating layer can detect conflicting state earlier and coordinate recovery. The critical case is that the same connectivity and reuse can enlarge the blast radius. The task is not to choose optimism or caution. It is to design evidence and authority so that useful adaptation does not become unbounded action.

Physical harm is not the only loss

Risk begins with physical injury but extends to quality, equipment, uptime, data, workforce and liability. A robot can place the wrong component without hitting anyone. A scheduling agent can create queues, starve a process or release work against an invalid recipe. A twin can support a confident decision after its state has drifted. A remote service can restore software while work in process remains physically inconsistent.

These outcomes have different time-to-harm and reversibility. A planning recommendation may permit review. Motion near a person may not. A wrong quality decision may remain latent until downstream assembly or field use. A cyber compromise may first look like a reliability problem. Governance therefore cannot be one universal “human in the loop.” It must match the decision, information, response time and consequence.

Workers also experience risk through surveillance, deskilling, work intensification and unclear responsibility. Consultation and training matter, but they are not substitutes for engineered protection. Nor does worker acceptance prove that a technical hazard is controlled. The report treats people as participants in design and operation, not as a flexi-

ble safety layer that compensates for missing controls.

Legal responsibility, privacy and product liability vary by application and jurisdiction. This report does not provide legal advice. Its narrower demand is traceability: who supplied the model, who integrated the application, who authorized the operating mode, who retained data, who changed the configuration, and who had authority to stop and restart.

Table 11. Risk, mechanism and accountable question

Risk domain	Mechanism to examine	Accountable question
Physical harm	Wrong state, motion or access during normal/degraded/nonroutine work	Which independent protective function limits the hazard?
Objective/uncertainty	Correct optimization of the wrong target; overconfident inference	When must the model abstain or be overridden?
Quality/process	Command succeeds but physical result is wrong	What direct readback establishes accepted work?
Cyber/common cause	Shared identity, update, data or platform failure reaches several cells	What remains locally valid and how is blast radius bounded?
Data/workforce	Excess collection, surveillance, skill loss or hidden supervision	What is necessary, permitted, contestable and owned?
Liability/change	Version or authority is unclear after modification	Who approved the change and which evidence was regenerated?
Recovery	Services return before physical state is reconciled	Who may reset, who may restart, and on what evidence?

Authors' governance table. It is not a risk assessment or allocation of legal liability.

Mechanism one: jam clearing changes the operating state

Imagine a cell stopping after inconsistent part-position and drive-status signals. Vision suggests a displaced part, but glare leaves open whether another object is present. The twin shows the preceding cycle because an event arrived late. The reasoning layer proposes a short reverse movement to expose the jam.

This is not a real incident. It is a mechanism for testing boundaries. The recommendation may be operationally sensible and still lack authority. Jam clearing has moved the application from routine production into nonroutine work. Maintenance or recovery state must dominate the throughput objective.

OSHA's current Control of Hazardous Energy overview states that servicing and maintenance can expose workers to unexpected energization, startup or release of stored energy and points to 29 CFR 1910.147. This is a U.S. regulatory baseline, not proof that a particular robot application complies.[45]

In the illustrative trace, authority logic withholds the proposed motion pending the site's validated access and hazardous-energy procedure. Basic

control remains in its defined non-production state. In parallel, the protective-function path continues observing the application-defined safety signals and can limit or stop motion regardless of the model recommendation.

The technician's physical inspection may reveal a geometry different from the prediction. After the work, recovery must reconcile equipment position, tooling, guards or protective devices, configuration and process state. A reset acknowledges or clears a defined condition. It is not the same as restart. Restart authorizes renewed operation after the required checks, by a site-defined role, in an application-defined mode.

A safety stop and an emergency stop are also not interchangeable. Their meaning and behavior belong to the application and applicable standards. Emergency stop is a complementary protective measure, not normal production control, the primary risk-reduction strategy or energy isolation.

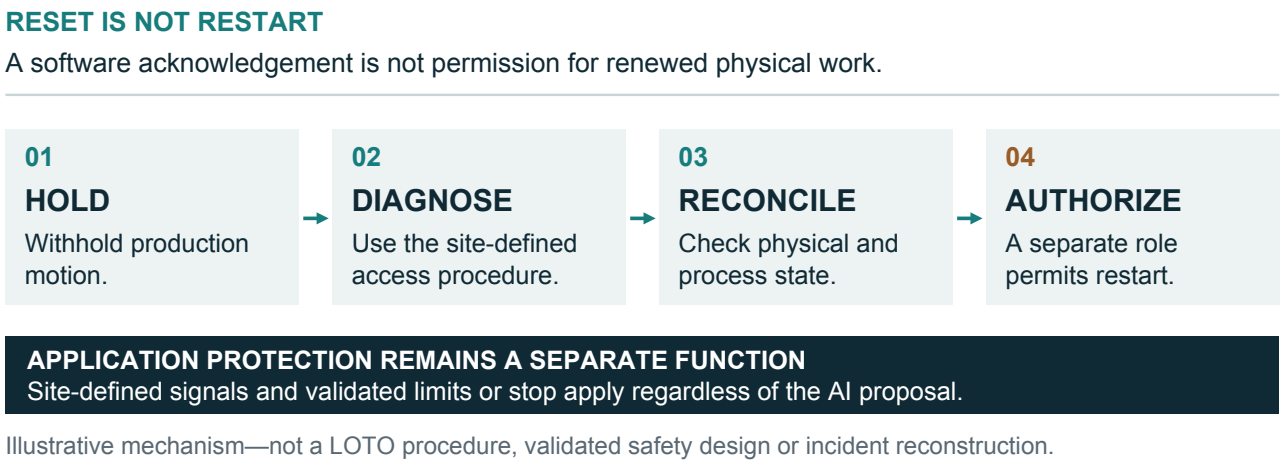


Figure 16. Nonroutine work and restart. Illustrative application: diagnosis and physical-state reconciliation precede a separately authorized restart.

Mechanism two: a plausible shared state can be wrong everywhere

Now imagine a shared platform coordinating several cells after an update. One connector repeats stale "ready" values while another loses authentication. The values remain syntactically valid, so a range check passes. The twin converges on an internally consistent but physically wrong state and recommends releasing work downstream.

The danger is not that a central model is always worse. Central context may improve scheduling and expose conflicts no local system sees. The danger is common cause: one wrong state can generate many locally plausible recommendations.

NIST SP 800-82 Rev. 3 describes operational technology in terms of interaction with the physical environment and treats security alongside performance, reliability and safety. It supports controls such as identity, segmentation, logging, availability, change management and recovery, but it does not prescribe the safe state of a specific plant.[46]

Command envelopes can reduce the blast radius: each local controller accepts only actions valid for its current mode and constraints. In a separate parallel path, application-validated protective functions continue to observe their defined local signals. They do not accept the shared platform's inference as a safety command.

Cross-cell anomaly detection compares platform state with local readback and identifies divergence. Unless independently safety-rated and validated, it is an operational diagnostic—not the protective function. The platform abstains from further coordination and declares degraded operation. The plant moves to a predeclared fallback such as local sequencing, fixed routing or controlled stop, with known capacity and sustainable duration.

Recovery does not end when services show green. Authorized versions and identities must be restored; work in process and equipment state must be reconciled; protective functions and command boundaries must be verified; restart authority must be explicit. Otherwise, a technically successful rollback can restart a physically inconsistent process.

NIST SP 1339 links OT backups with change management, regular restoration tests and recovery exercises. It calls for the engineering software, compatible hardware and documentation needed to restore operations.[47] For a plant, technical restoration must be followed by the checks that establish the physical process is ready to resume.

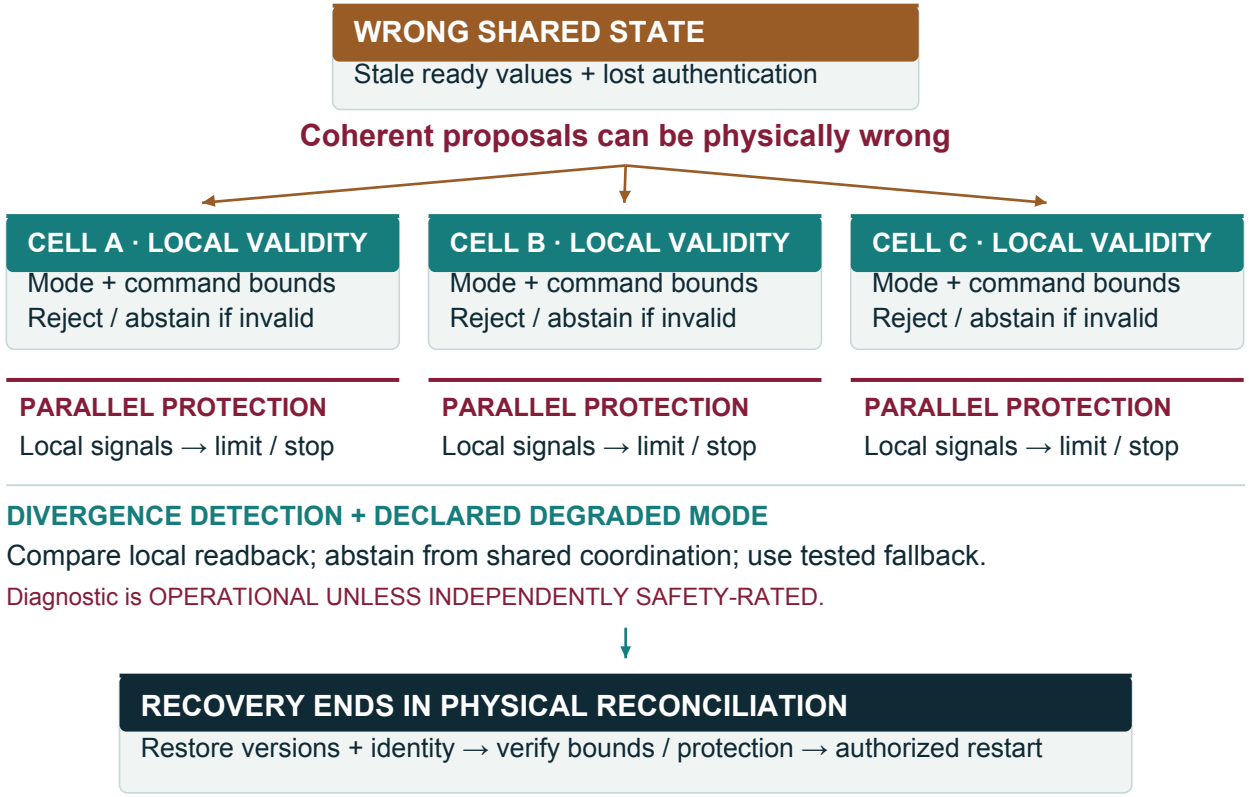


Figure 17. Shared error and local containment. Illustrative application: a shared-state error reaches several assets. Local validity checks, protective functions and a defined fallback limit the consequences.

U.S. robot safety and application responsibilities

ANSI/A3 R15.06-2025 is the U.S. national adoption of ISO 10218-1:2025 and ISO 10218-2:2025. Its published scope distinguishes industrial robots from robot applications and cells. ANSI/A3 R15.06-3-2025 addresses use of industrial robot cells.[48] The ISO publications provide the corresponding international framework.[49, 50] A project must establish the applicable requirements with qualified safety and integration specialists; the product name alone establishes no application-level compliance.

That distinction prevents a common category error. A robot product can include safety-related capabilities while the completed application remains unsafe or unvalidated. Tool, payload, layout, speed, modes, access, neighboring equipment and recovery all matter. A “cobot” label does not certify the work cell. A Digital Twin cannot validate the safety of the same action it recommends merely by reporting that its internal state is consistent.

The architecture must also avoid false independence. If operational inference and protective monitoring depend on the same stale signal, identity service or update path, two boxes in a diagram do

not create two defenses. The application must establish the required degree of functional independence and control common-cause failure. That is an engineering claim to validate, not a phrase to assume.

Wrong objective, uncertain world

Some failures occur even when sensing and control perform as designed. A scheduler minimizing delay can increase risky manual intervention. A quality optimizer can reject too much product. A maintenance model can generate so many alerts that attention becomes the scarce resource. These are objective and organizational failures, not merely model-accuracy problems.

Uncertainty must reach the decision. A confidence score is useful only if it changes behavior: seek another observation, narrow authority, route to a person, use a baseline or stop the operation. The relevant test is not whether a model is uncertain in general. It is whether uncertainty in the state, parameter or measurement propagates into an acceptance rule appropriate for the consequence.

Human supervision works best when the action is slow enough, reversible enough and legible enough for intervention. It fails as a control argument when the person lacks time, context, competence or real authority. Required intervention time and available intervention time are measurements for the application. There is no universal number to insert into this report.

Recovery is a product capability

Normal-cycle performance attracts attention; recovery determines trust. A buyer should ask who diagnoses the fault, how remote work is recorded, what the local team can do without the supplier, how versions are restored, and what happens when the vendor or network is unavailable. A supplier should count that labor rather than hiding it behind an uptime headline.

The stop rule is straightforward even when thresholds are not: do not extend physical authority where ownership, protection, readback or recovery is undefined. Continue a pilot only when the next uncertainty is measurable and the fallback is acceptable. Expand only when another operating context and its correlated failure paths have been tested.

A system that handles variation, recognizes when to defer a decision, contains failure and returns to controlled operation offers a credible route to wider industrial use. Recovery belongs in the product specification because it determines how much useful autonomy the plant can sustain.

10 / OUTLOOK

Outlook

The next decade is likely to bring more capable models into factories that retain diverse equipment, processes and responsibilities. Our outlook favors a heterogeneous manufacturing system, with selective learning-based autonomy and growing demand for coordination. The pace will depend on measurement, transfer between tasks, lifecycle economics and industrial governance. The reference date is October 2026.

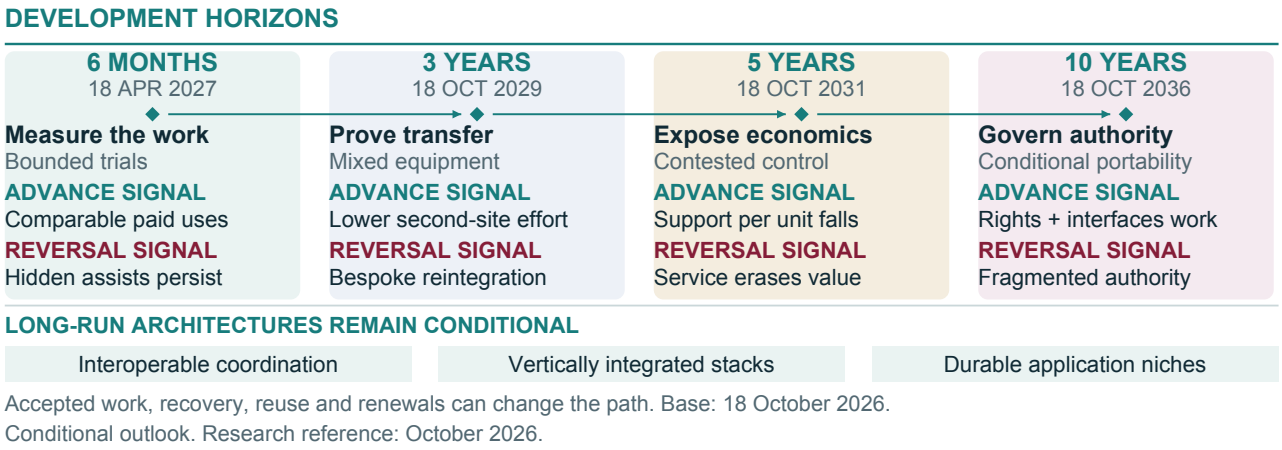


Figure 18. Development horizons and evidence. The outlook tracks measurement, transfer, lifecycle economics and governance. Changes in operating evidence can accelerate or reverse the pathways.

Six months: measurement is the bottleneck

Over the next six months, the central pathway is still bounded trials, assistance and integration. Six months is enough to improve test design and expose hidden labor; it is rarely enough to establish lifecycle economics.

Sergey Levine's warning is useful: "The scope will have to start out small." In context, he describes capability expanding gradually with human participation; this is a founder-researcher's pathway, not a factory forecast.[51]

The decisive information is denominator discipline: scheduled hours, all attempts, accepted output, human assists and recovery. Several repeat paid uses with comparable boundaries would move priority toward replication. Continued dependence on undisclosed teleoperation, subsidy or exceptional engineering would move it back toward instrumentation and baseline repair.

Our near-term judgment favors suppliers that make assistance and support requirements visible and customers that establish a usable baseline before negotiating replication.

Three years: transfer becomes the bottleneck

Over three years, the central pathway is a heterogeneous factory: dedicated cells, industrial arms, mobile robots, human work and learning-based systems coexist. That is more plausible than rapid convergence on a universal machine because installed equipment turns over slowly and useful tasks impose different physical constraints.

The question changes from "did one deployment work?" to "what was reused?" The relevant evidence is the engineering, validation and support required for a second task or site. Portable interfaces, task representations and recovery procedures would shift priority toward broader platforms. If every installation remains a bespoke integration

project, incumbents and integrators retain the advantage even as models improve.

The coordination thesis becomes more important as heterogeneous assets encounter shared constraints. Evidence of economical reuse across tasks would support investment in a shared operating context. A cohort of general-purpose machines transferring with little intervention would strengthen a different route: broader capability concentrated at the endpoint.

Five years: lifecycle economics reveals the control point

Over five years, published comparable cohorts—if suppliers and customers produce them—may allow architectures to be judged over renewal, support and second-site deployment. The passage of five years alone proves nothing. Without cohorts, the control point remains a hypothesis.

The central pathway is contested control. OEMs may win where tight endpoint responsibility dominates; industrial software incumbents where semantics and service relationships matter; integrators where brownfield adaptation remains bespoke; independent operating layers where cross-vendor coordination creates a repeatable gain. The deciding evidence is contribution after field engineering, support, warranty and financing, together with portability and recovery when a shared layer is unavailable.

This is the horizon at which architecture must meet cash. Falling second-site effort and support per accepted unit would favor a reusable layer. Persistent reintegration, diffuse accountability or cheaper MES/MOM extensions would favor the incumbent route. Manufacturers should negotiate fallback and exit before dependency grows; suppliers should state whether they sell product, managed outcome or project engineering; investors should compare cohorts rather than recurring-revenue labels.

Ten years: governance limits the feasible market order

Over ten years, technical capability may no longer be the only rate limiter. Liability, command authority, interoperability, data rights and local operating competence can determine which architectures are acceptable. Three orders remain plausible: interoperable intelligence coordinating specialized machines; vertically integrated stacks owning interfaces and service; or durable fragmentation into application niches.

The central analytical view is conditional interoperability: common semantics, decision contracts and outcome records across specialized assets, with inference and protection remaining distributed by timescale and function. It does not require one central OS or model. It loses if vertical integration produces materially lower lifecycle burden, or if physical variance and responsibility prevent reusable coordination from earning its extra dependency.

Humanoid form may win tasks where mobility, reach and reuse outweigh complexity. It may lose to a fixture, arm, mobile base, person or redesigned process. A forecast of plant architecture cannot be inferred from a forecast of one embodiment.

What changes the path?

Revise upward if independent multi-site cohorts show economical transfer, sustained accepted output and declining support burden. Revise downward if customers stop deployments, recovery remains specialist-dependent, or service and financing consume the value attributed to flexibility. Where comparable public operating records are absent, the outlook remains conditional. Contract, operation, renewal and repeat site are different stages.

11 / BUSINESS DECISIONS: BUY THE NEXT JUSTIFIED STEP

Business decisions: buy the next justified step

The first 90 days should turn interest in Physical AI into a specific investment decision. Select one production problem, establish its cost, compare the credible alternatives and obtain the measurement most likely to change the choice. The plan below assigns that work to accountable functions and leads to a decision to buy, build, partner, continue testing or defer.

The first 90 days

Table 12. A 90-day evidence plan

Period	Decision work	Accountable function	Dependency and measurement
Days 1–15	Define the problem, accepted output and baseline; observe normal and exception work.	Operations owner with quality and finance	Stable task boundary; accepted units, demand, interventions, scrap/rework and downtime.
Days 16–30	Compare simpler alternatives and design a bounded test. Define data rights, authority, interruption and recovery.	Process/controls lead; EHS and IT/OT according to impact; procurement/legal for rights	Same-task comparator; approved test boundary; named stop condition.
Days 31–60	Run only the authorized test. Log every attempt, rejection, assist, fault, recovery and configuration change.	Pilot owner and integrator; quality owns acceptance	Traceable event log linked to accepted output; no inference from selected demo clips.
Days 61–75	Reconcile operational result with full buyer cost and demand. Test downside and missing fields.	Finance with operations and maintenance/service	Integration, ramp, supervision, support, validation, downtime, financing and exit counted once.
Days 76–90	Decide stop, continue, defer or conditional expand. Record owner, evidence and next gate.	Plant steering group; executive sponsor accountable	Predeclared thresholds; second-context plan and service capacity before expand.

Manufacturer: protect the task boundary

Define accepted work and the demand condition that gives it value. Compare process redesign, human work, local automation, an incumbent extension and a learning-based option on the same boundary.

Buy when a bounded solution has demonstrated the required result and lifecycle burden. Build only when the capability is strategically differentiating and the plant can own data, integration, validation and recovery. Partner when competence is external but responsibility remains explicit. Defer when one obtainable measurement could change the choice. Stop physical execution when authority, protection or recovery is undefined.

Expand means the result survives another context with declared support, version and fallback requirements.

Technology company or integrator: productize the exception

The deliverable is accepted work inside a declared operating envelope. State what the system sees, may decide and command, when it abstains and who restores operation.

Buy commodity components where differentiation is weak and rights are clear. Build task knowledge, interfaces and evaluation assets that reduce second-deployment effort. Partner for controls, safety engineering, field service or channels. Defer when buyer access or acceptance criteria are absent. Stop claims that require undisclosed human labor, unsupported safety or unmeasured ROI.

The commercial test is contribution after deployment, support, warranty, financing and recovery—not recurring revenue alone.

Investor: operating cohorts and cash

Treat contract, backlog, installed unit, accepted operation, renewal and repeat site as different evidence. Ask who owns command permission, process semantics, recovery and working-capital burden.

Continue diligence when a company can show bounded paid operation, cohort support and the path from first to second deployment. Defer when product margin excludes material field labor or when customer concentration substitutes for repeatability. Stop an investment thesis that depends on a category TAM assembled from overlapping robot, AI and manufacturing markets. Expand capital only when milestones improve both operating evidence and cash quality. This is a diligence framework, not advice regarding any security.

Every decision record should identify what is known, what remains to be measured, who owns the next action and when the decision will be reopened. Continuing a pilot requires a defined learning objective and an end point. Expansion requires a specified next task or site, support capacity and conditions for reversal.

THREE AUDIENCES. THREE DECISIVE TESTS.

MANUFACTURER	**Accepted output and full lifecycle burden.** Does the result survive the next operating context?
TECHNOLOGY COMPANY	**Repeatable delivery, exception handling and support.** What must be rebuilt or serviced at the second site?
INVESTOR	**Operating cohorts, contribution and cash quality.** Is deployment becoming repeatable—or merely recurring?

Decision lens. Three audiences evaluate different outcomes; selected demonstrations cannot replace an operating record.

A filled decision record

For the inspection conflict in Chapter 5, the immediate issue is the unvalidated estimate of B's duration under a changed part condition. The decision record below turns that uncertainty into a task for operations and quality, while preserving the approved production route.

Table 13. A filled decision record

Field	Recorded decision
Decision now	**Defer** changing the approved sequence; **continue** only a shadow timing test. The incumbent rule wins while it produces the feasible answer at lower burden. If the changed condition invalidates that baseline, **stop** execution pending review.
Missing field	Whether a changed part condition makes B take five minutes. The estimate is unvalidated.
Owner	Operations owns timing and the sequence; quality owns acceptance of the changed condition. Neither a model nor supplier can silently redefine expiry or approve output.
Next test	Observe B under the declared changed condition without issuing commands. Compare measured completion, accepted order outcome and exception work with the incumbent estimate.
Reopen when	Operations and quality validate or reject the five-minute estimate. If no feasible sequence remains, **stop** and escalate the constraint; do not call an infeasible answer "optimized."
Expansion gate	Only after the new route improves a repeated cross-asset decision over the incumbent, within authority, recovery and lifecycle-cost limits.

This record connects evidence to authority and spending. It also prevents repeated data collection that cannot change the decision. The most valuable next measurement is the one that distinguishes the credible alternatives and leads to a named action.

DBR77 RESEARCH / TECHNICAL NOTE

Technical note

The example in Chapter 6 is a planning model with synthetic inputs. Its values are normalized, not dollar prices or observed customer returns. It compares one local manipulation task. It does not estimate the value of a plant-wide coordination system.

Output and accepted work

H = 1,000 available time units per period. q = 0.95 is the share of all attempts producing an accepted result. c is the time per attempt excluding additional recovery. e is the share of the same attempts needing recovery. r is additional time per recovery event.

$G = Hq / (c + er)$.

This average-rate approximation assumes stable times and no overlapping recovery. It omits queues, correlated failures and rare extended outages. The baseline produces $G_0 = 950$ accepted units per period.

Demand and value

D caps additional usable demand per period. m = 1 is contribution per additional accepted unit after unit-variable costs. $\Delta G = G - G_0$.

$B = m[\min(\max(\Delta G, 0), D) - \max(-\Delta G, 0)]$.

$F = B - C$. $V = 3F - K - X$.

The benefit formula caps positive contribution at usable demand and preserves lost contribution below baseline. C is incremental recurring cost, K initial cost and X exit cost. V covers three identical periods without discounting, tax, financing or residual value. It is not NPV. Recovery capacity loss and separate service expense must be counted without duplicating the same financial loss.

Sensitivity

Under the base assumptions, O_3 exceeds O_2 above approximately 148.75 additional demand units. At D = 220, their values are equal at an O_3 recovery-event share of approximately 0.02154. These unrounded model thresholds explain the rounded business discussion. They are economic comparisons, not acceptable safety error rates.

The separate supplier illustration has revenue of 100, operating delivery costs of 85 and contribution of 15 before R&D, overhead and tax. Initial funding of 60 less 20 collected before operation creates a funding gap of 40. That gap is a financing requirement, not an extra expense deducted from the same contribution.

Applying the model

Replace synthetic inputs with a plant-approved record of all attempts, accepted output, demand and costs. Use period-specific ramp and working capital for a real proposal. Compare the same orders and quality requirements across alternatives. A range spanning different preferred options identifies the measurement to obtain before committing investment.

DBR77 RESEARCH / SOURCES & NOTES

Sources & Notes

Numbers identify the publications cited in the text. Repeated uses of the same publication share an entry. URLs are printed and linked for digital and paper reading.

1. NVIDIA, “What is Physical AI?”, institutional glossary, undated, accessed 18 October 2026; definition and examples. www.nvidia.com/en-gb/glossary/generative-physical-ai/

2. Fei-Fei Li, “Fei-Fei Li on Spatial Intelligence and Human-Centered AI,” Possible publisher transcript; discussion of large world models and LLMs. Undated; AI-assisted, lightly edited transcript. www.possible.fm/podcasts/feifei/

3. Google DeepMind, “Gemini Robotics brings AI into the physical world,” 12 March 2025; model descriptions. deepmind.google/blog/gemini-robotics-brings-ai-into-the-physical-world/

4. John Leonard, “Byte Bites Transcript,” MIT CSAIL, 14 May 2024, discussion of representations and robot decisions. cap.csail.mit.edu/john-leonard-byte-bites-transcript

5. Yevgen Chebotar and Tianhe Yu, “RT-2: New Model Translates Vision and Language into Action,” Google DeepMind, 28 July 2023. deepmind.google/blog/rt-2-new-model-translates-vision-and-language-into-action/

6. Dibya Ghosh et al., “Octo: An Open-Source Generalist Robot Policy,” Robotics: Science and Systems, 2024, section IV, Table I and section V. www.roboticsproceedings.org/rss20/p090.pdf

7. NVIDIA GEAR, “GR00T N1.6: An Improved Open Foundation Model for Generalist Humanoid Robots,” 15 December 2025. research.nvidia.com/labs/gear/gr00t-n1_6/

8. Emma Brunskill, “Lecture 1: Introduction to RL,” Stanford CS234, Winter 2026, slides on reinforcement learning, optimization and sequential decisions. web.stanford.edu/class/cs234/slides/lecture1post.pdf

9. Jack Parker-Holder and Shlomi Fruchter, “Genie 3: A New Frontier for World Models,” Google DeepMind, 5 August 2025. deepmind.google/blog/genie-3-a-new-frontier-for-world-models/

10. Yann LeCun, Lex Fridman Podcast #416, 7 March 2024, approximately 01:19:42–01:28:54, publisher transcript. lexfridman.com/yann-lecun-3-transcript

11. Guodong Shao, Joe Hightower and William Schindel, “Credibility Consideration for Digital Twins in Manufacturing,” Manufacturing Letters 35, pp. 24–28, 2023; NIST author manuscript, sections 1–3. www.nist.gov/publications/credibility-consideration-digital-twins-manufacturing

12. Wilbert Pumacay et al., “COLOSSEUM: A Benchmark for Evaluating Generalization for Robotic Manipulation,” Robotics: Science and Systems, 2024. www.roboticsproceedings.org/rss20/p133.html

13. Maximilian Metzner, “Advanced Robotics and AI in Industrial Manufacturing — Episode 1,” Siemens, 6 June 2025, publisher transcript. blogs.sw.siemens.com/thought-leadership/advanced-robotics-and-ai-in-industrial-manufacturing-episode-1-transcript/

14. Rainer Brehm and David Humphrey, “Industrial AI Is Already on the Shop Floor,” Siemens Future Ready, 9 September 2026, publisher transcript, 03:25 and 07:46–10:49. blogs.sw.siemens.com/thought-leadership/industrial-ai-is-already-on-the-shop-floor-transcript/

15. BMW Group, “BMW Group to Deploy Humanoid Robots in Production in Germany for the First Time,” 27 February 2026, Spartanburg section. www.press.bmwgroup.com/united-kingdom/article/detail/T0455870EN_GB/bmw-group-to-deploy-humanoid-robots-in-production-in-germany-for-the-first-time

16. BMW Group, “BMW Group Advances the Use of Physical AI in Production with Figure 03 Project in Spartanburg,” 25 June 2026. www.press.bmwgroup.com/global/article/detail/T0458778EN/bmw-group-advances-the-use-of-physical-ai-in-production-with-figure-03-project-in-spartanburg

17. Mercedes-Benz Group, “AI and Humanoid Robots,” 18 March 2025, Apollo section. group.mercedes-benz.com/company/production/production-network/mbdfc-humanoid-robots.html

18. GXO Logistics, “GXO Signs Industry-First Multi-Year Agreement with Agility Robotics,” 27 June 2024. investors.gxo.com/news-releases/news-release-details/gxo-signs-industry-first-multi-year-agreement-agility-robotics/

19. BMW Group, “Intelligently Connected Factory: BMW Group Plant Regensburg Drives Digitalisation and Automation of Logistics with Cloud-Based Traffic Control System,” 29 October 2024, pp. 1–4. www.press.bmwgroup.com/belux/article/attachment/T0446014FR/623985

20. Universal Robots, “Precise and Cost-Effective UR Cobots Enable Increased Output,” RCM Industries customer case, undated. www.universal-robots.com/case-stories/rcm-industries/

21. FANUC America, “Closure Systems International Automates Material Handling to Boost OEE,” 15 November 2022, updated 21 July 2026. www.fanucamerica.com/case-studies/closure-systems-international

22. International Federation of Robotics, “Robot Density Surges in Europe, Asia, and Americas,” 8 April 2026, United States and 2024 reference data. ifr.org/ifr-press-releases/news/robot-density-surges-in-europe-asia-and-americas

23. Bonney et al., The Microstructure of AI Diffusion: Evidence From Firms, Business Functions, and Worker Tasks, CES-WP-26-25, April 2026, Table 1(a), Panel A, and sections 2.1 and 3.1. Firm-weighted estimate: 17.7%; the abstract rounds this to 18%. www2.census.gov/ces/wp/2026/CES-WP-26-25.pdf

24. U.S. Bureau of Labor Statistics, Job Openings and Labor Turnover — August 2026, 29 September 2026, Table 1, manufacturing row, seasonally adjusted and preliminary. www.bls.gov/news.release/archives/jolts_09292026.pdf

25. Board of Governors of the Federal Reserve System, *Industrial Production and Capacity Utilization — September 2026*, G.17 Table 7, released 16 October 2026. www.federalreserve.gov/releases/g17/20261016/g17.pdf

26. NIST, Annual Report on the U.S. Manufacturing Economy: 2025, NIST AMS 100-76, February 2026, printed pp. 50–55, discussion of competitiveness and reshoring. nvlpubs.nist.gov/nistpubs/ams/NIST.AMS.100-76.pdf

27. U.S. Department of Commerce, "Restoring American Semiconductor Manufacturing Leadership Through an Agreement on Trade & Investment with Taiwan," 15 January 2026. www.commerce.gov/news/fact-sheets/2026/01/fact-sheet-restoring-american-semiconductor-manufacturing-leadership

28. Manufacturing.gov, "Advanced Manufacturing," federal announcement of 22 July 2026, NIST funding opportunity for 14 MEP centers. www.manufacturing.gov/topic/advanced-manufacturing

29. Bryan DeBois, *The Future of U.S. Manufacturing: Leveraging AI and Smart Factories*, U.S. Manufacturing Today / Veryable, n.d., publisher transcript at 03:03–05:31 and 28:00–31:00. www.veryableops.com/u-s-manufacturing-today-podcast/episode-30

30. NIST MEP, "AMG Industries Increased Productivity and Cost Savings with Collaborative Robots," 31 July 2024. Program/customer-reported results. www.nist.gov/mep/successstories/2024/amg-industries-increased-productivity-and-cost-savings-collaborative-robots

31. NIST MEP, "JCGibbons Manufacturing: Integrating Cobots for Improved Efficiencies," 28 October 2022, updated 21 June 2023. Program/customer-reported results. www.nist.gov/mep/successstories/2022/jcgibbons-manufacturing-integrating-cobots-improved-efficiencies

32. Brian Weinstein, "Meet the Manufacturers: Chapco," ManufactureCT, undated, publisher transcript, 10:15–11:33. manufacturect.org/podcast/chapco/

33. Julian Counihan, "Why You Can't Vibe-Code a Factory," Alpha Partners, 1 July 2026, publisher transcript. alphapartners.com/podcast/julian-counihan-schematic-ventures/

34. Symbotic Inc., FY2025 Form 10-K, filed 24 November 2025, Business, customer concentration and remaining-performance-obligation sections. www.sec.gov/Archives/edgar/data/1837240/000183724025000278/sym-20250927.htm

35. Symbotic Inc., FY2025 investor presentation filed with the SEC, "Symbotic at a Glance." www.sec.gov/Archives/edgar/data/1837240/000183724025000277/symboticinvestorpresenta.htm

36. Teradyne, FY2025 Form 10-K, Robotics, MD&A, inventory and Note Q. investors.teradyne.com/sec-filings/all-sec-filings/content/0001193125-26-059002/ter-20251231.htm

37. Rockwell Automation, FY2025 Annual Report, segment descriptions, results and revenue-recognition note. www.rockwellautomation.com/content/dam/rockwell-automation/documents/pdf/company/about-us/ir/2025/ROK-AR-2025-Final.pdf

38. Siemens AG, Annual Report 2025, Digital Industries strategy and segment results. assets.new.siemens.com/siemens/assets/api/uuid%3A428ea18a-e7ab-4f93-a160-33908f1c3540/Siemens-Annual-Report-2025.pdf

39. Darragh de Stonndún, *Automation as a Present-Day Necessity: Building Scalable Robotics Platforms*, U.S. Manufacturing Today / Veryable, n.d., publisher transcript at 04:32, 18:11–20:55 and 33:51–34:50. www.veryableops.com/u-s-manufacturing-today-podcast/episode-53

40. Jeff Wilke, *Deployed: Where Physical AI Gets Real — Scale like Amazon: How Jeff Wilke is Rebuilding USA Manufacturing*, Cobot, published 24 March 2026, publisher transcript, discussion of retained technical competence. share.transistor.fm/s/b27cbf3a

41. Daron Acemoglu and Pascual Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy* 128(6), electronically published 22 April 2020; pp. 1, 13 and 54. Author-hosted publisher PDF economics.mit.edu/sites/default/files/publications/Robots%20and%20Jobs%20-%20Evidence%20from%20US%20Labor%20Markets.p.pdf

42. Georg Graetz and Guy Michaels, *Robots at Work*, author-hosted version dated 14 February 2018, pp. 1–8. Author-hosted PDF personal.lse.ac.uk/michaels/graetz_michaels_robots.pdf

43. David Binder Research for AFL-CIO, *AI and Work Survey*, fielded 14–22 April 2026; memo n.d., both pages including methodology. aflcio.org/sites/default/files/2026-05/DBR_AFLCIO_AI_Research_Memo.pdf

44. Liz Shuler, AFL-CIO President, National Press Club speech, 28 April 2026, published speech text. aflcio.org/speeches/afl-cio-president-liz-shulers-speech-national-press-club

45. U.S. Occupational Safety and Health Administration, "Control of Hazardous Energy (Lockout/Tagout)," current overview accessed 18 October 2026. www.osha.gov/control-hazardous-energy

46. NIST SP 800-82 Rev. 3, *Guide to Operational Technology (OT) Security*, final 28 September 2023. Selected full-text sections: §§3.1, 4.2.2, 5.2.3.1–2 (segmentation/logging), 5.3.1–2 (safety/availability), 6.5 (recovery), F.7.5/CM-3 (change) and F.7.7 (identity). Publication record csrc.nist.gov/pubs/sp/800/82/r3/final; official PDF nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.pdf

47. NIST SP 1339, OT Backup Quick Start Guide, June 2026, pp. 1–2, backup planning, restoration testing and engineering documentation. nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1339.pdf

48. ANSI/A3 R15.06-2025, Parts 1 and 2, and ANSI/A3 R15.06-3-2025, Use of Industrial Robot Cells. Public scope and national-adoption information. webstore.ansi.org/standards/ria/ANSIA3R15062025PartsPart

49. ISO 10218-1:2025, Robotics — Safety Requirements — Part 1: Industrial Robots, February 2025. Public catalog scope. www.iso.org/standard/73933.html

50. ISO 10218-2:2025, Robotics — Safety Requirements — Part 2: Industrial Robot Applications and Robot Cells, February 2025. Public catalog scope. www.iso.org/standard/73934.html

51. Sergey Levine, "Fully autonomous robots are much closer than you think," *Dwarkesh Podcast*, 12 September 2025, publisher transcript, "Timeline to widely deployed autonomous robots." www.dwarkesh.com/p/sergey-levine

52. DBR77, company and product descriptions, accessed 18 October 2026. Company dbr77.com/en/, Vector vector.dbr77.com/en, IRIS iris.dbr77.com/, Digital Twin dt.dbr77.com/en, Marketplace marketplace.dbr77.com/marketplace

DBR77 RESEARCH / ABOUT DBR77

About DBR77



DBR77 develops software and hardware for industrial digital transformation. Its portfolio includes industrial IoT, Digital Twin, IRIS and an automation Marketplace, alongside Vector, its industrial language-model proposition.[52]

Our strategic direction is to build an autonomous Plant Operation System: an operating intelligence that connects perception, process learning, simulation and human operational knowledge with bounded decisions and verified results. The ambition is coordination across heterogeneous equipment, robots and the work of people, with explicit authority and safeguards for each application.

Different parts of the ecosystem serve different purposes. Perception helps establish the state of work. Simulation tests alternatives within a model. Synthetic data supports selected learning and evaluation tasks, followed by real-world validation. The Marketplace connects production challenges with suppliers and automation capabilities. Together, these activities inform our approach to plant-level operating intelligence.

This is the perspective from which DBR77 contributes to the Physical AI debate. The report evaluates this development direction against the same questions of evidence, alternatives, lifecycle cost and responsibility applied to other approaches. Demonstrated capabilities and future ambitions remain distinct.

dbr77.com · contact@dbr77.com

Explore the ecosystem

AI Vision · Industrial IoT · Plant Operation System · Digital Twin · Marketplace · Consultify

dbr77.com

Our development agenda

Connect useful perception with process models and operational knowledge. Test decisions in simulation. Verify consequences in the plant. Reuse the resulting knowledge across changing tasks and heterogeneous equipment. This is the direction in which DBR77 is developing its autonomous Plant Operation System.

DBR77

INDUSTRIAL INTELLIGENCE

Physical AI in American Manufacturing



An industrial perspective on machine intelligence, plant coordination and the next investment decision.

DBR77 Robotics Inc. / dbr77.com

AI-generated conceptual illustration

IEX 2026