DBJ Research

No. 445 (August 7, 2026)

Social Implementation of Physical AI:
Industrial Structure and Policy for Delivering Technology to the Field

Industrial & Regional Research Dept. Hiromu Samuta (currently Niigata Branch), Haruto Suzuki

Executive Summary

- The generative AI market's competitive axis is shifting from the scale of general-purpose models toward domain-specific AI. Competitiveness now hinges on "integrated implementation capability"—unifying field data, hardware, and operational know-how—especially where AI operates in physical space: the domain of physical AI.
- Japan's strengths—accumulated field data, a world-class industrial-robot base, and long-standing operational know-how—remain underused amid a shortage of systems integrator ("SIer") functions for implementation and integration. The robot SIer industry is structurally difficult both to scale and to fund.
- This paper examines joint ventures between AI vendors and user companies as a framework to fill this gap and foster SIers, and proposes demand-side support for physical AI—demonstration subsidies and tax incentives—conditioned on domestic data management.

1. Shifts in the Competitive Axis of the AI Market

(1) Trends in the Generative AI Model Market

Innovation centered on generative AI is advancing worldwide, and the economy and society are entering a new phase. It is driven by large language models (LLMs), which learn from vast web text to gain advanced language understanding and generation. LLMs now underpin text generation, summarization, and translation, spurring diverse applications. They also substitute for and complement part of the intellectual activity humans perform through language (Figure 1-1).

Figure 1-1 Generative AI structure and use process

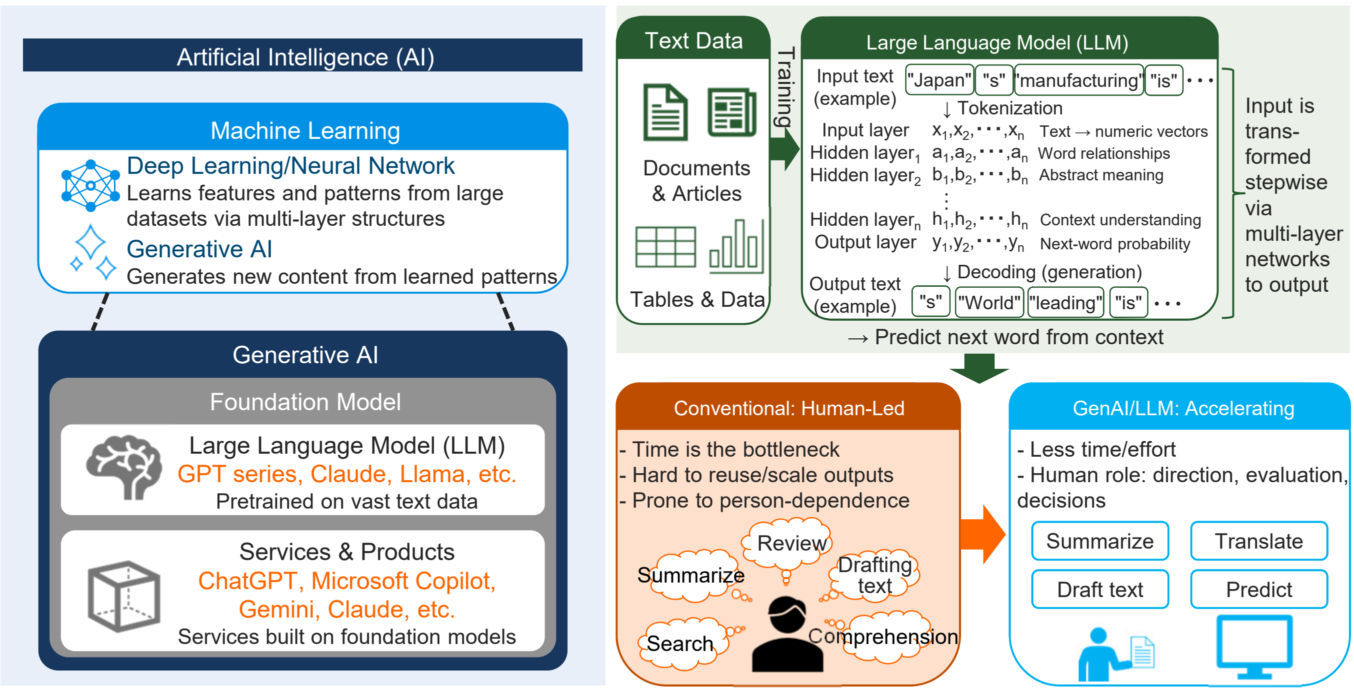

(Note) Prepared by Development Bank of Japan Inc. from various sources.

(2) The Expanding Domain-Specific Language Model Market

The generative AI model market has entered a phase of rapid expansion. According to a March 2026 Gartner, Inc. forecast, end-user spending will grow from about $15.4 billion in 2025 to about $152.5 billion by 2030, a compound annual growth rate (CAGR) of 58.2% (constant currency, 2024 base) (Figure 1-2). Most of this market is still held by foundation (general-purpose) models, but domain-specific language models (DSLMs) for sectors such as finance, healthcare, and manufacturing are gaining ground. The DSLM market is expected to grow from about $1.4 billion in 2025 to about $18.9 billion in 2030, a CAGR of 68.7% that outpaces the overall market's 58.2% (Figure 1-2). DSLMs tend to be more accurate and cheaper to run than general-purpose models, easing adoption in real operations. Gartner estimates "By 2028, more than 60% of GenAI models used by enterprises will be domain-specific — a steep rise from less than 10% in 2024."[1] Lower adoption barriers from open-source diffusion and more efficient training, plus early uptake in regulated industries, should expand the ecosystem of data platforms and inference infrastructure. That should lift demand for IT services such as application implementation, infrastructure, and consulting, in line with the growth outlook for the overall IT services market (Figure 1-3).

[1] Gartner®, Top Strategic Technology Trends for 2026: Domain-Specific Language Models, Samantha Searle et al., 18 October 2025 GARTNER is a trademark of Gartner, Inc. and its affiliates.

Figure 1-2 Generative AI model market end-user spending (2024-30, calendar year)

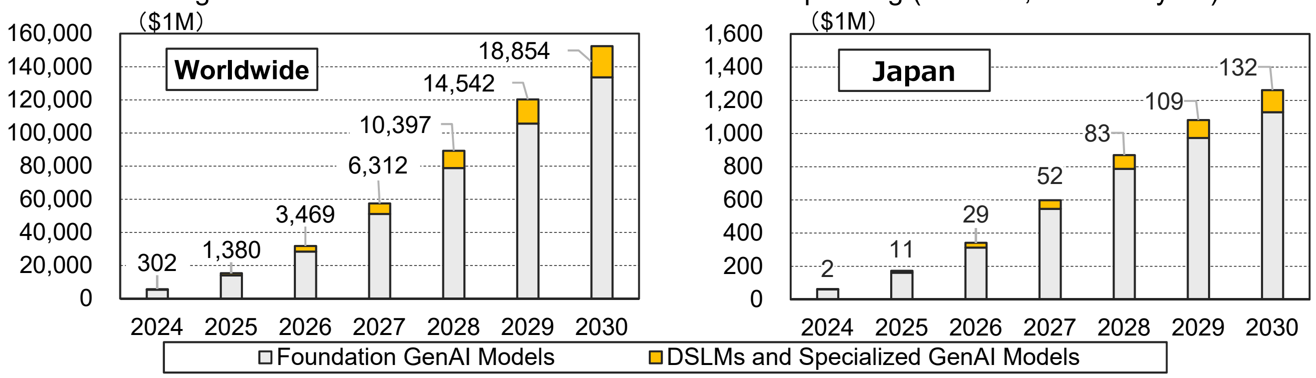

(Notes) Graphs created by DBJ based on Gartner research. Calculations performed by DBJ. Gartner®, Forecast: Generative AI Models, Worldwide, 2024–2030, 1Q26, Arunasree Cheparthi et al. 25 March 2026. End-User Spending in Constant Currency basis. GARTNER is a trademark of Gartner, Inc. and its affiliates. Constant currency basis (2024 base; USD 1 = JPY 151.47).

Figure 1-3 Global IT services market size and growth outlook (2024-30, calendar year)

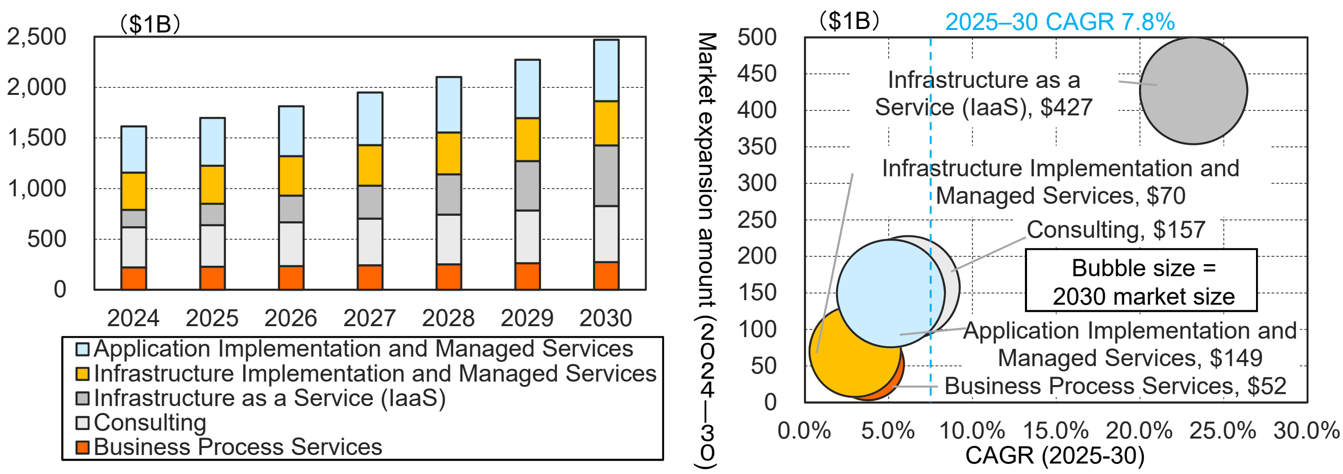

(Notes) Graphs created by DBJ based on Gartner research. Calculations performed by DBJ. Gartner®, Forecast: Services, Worldwide, 2024–2030, 1Q26, Neha Sethi et al., 25 March 2026. End-User Spending Constant Currency basis. GARTNER is a trademark of Gartner, Inc. and its affiliates. Constant currency basis (2024 base; USD 1 = JPY 151.47).

(3) Areas Where Japan Should Create Value: Field data, hardware integration, operational know-how

Given the shift "from general-purpose to domain-specific," where does Japanese industry's advantage lie? General-purpose LLMs demand massive compute and investment; so, the race is led by hyperscalers (Figure 1-4). These firms invest tens of trillions of yen a year in data centers and AI infrastructure; by 2026 major U.S. tech firms' combined spend should top ¥100tn, rivaling Japan's general-account budget. Japan is thus unlikely to win the "scale" race for general models head-on.

As competition shifts to domain-specific AI, Japan can draw on advantages absent in the general-model race, with strengths in three areas: (1) accumulated field data with access restricted by regulation; (2) a concentrated hardware base as the world's largest maker of industrial robots; and (3) implementation and operational know-how honed over years (Figure 1-5). Rather than the "scale" race, Japan should compete on "integrated implementation capability"—uniting field data, hardware, and operational know-how with AI models and raising value through continuous improvement, above all in physical AI.

Figure 1-4 Major U.S. tech firms' capex vs. Japan's general account

Company / Government	2025 actual	2026 guidance
Amazon	$131.8B (¥20.4tn)	~$200B (¥31.0tn)
Microsoft	$83.1B (¥12.9tn)[1]	~$190B (¥29.5tn)[1]
Alphabet (Google)	$91.4B (¥14.2tn)	$175–185B (¥27.1–28.7tn)
Meta	$72.2B (¥11.2tn)	$125–145B (¥19.4–22.5tn)[2]
Apple	$12.7B (¥2.0tn)[3]	~$14B (¥2.2tn)[4]
GAFAM 5 total	$391.2B (¥60.6tn)	~$704–734B (¥109.1–113.8tn)
Japanese govt general account	$745B (¥115.5tn)	$789B (¥122.3tn)

(Notes) Prepared by Development Bank of Japan Inc. from various sources. USD 1 = JPY 155
[1] Microsoft's fiscal year is Jul–Jun. 2025 actual $83.1B is calendar-year basis (quarterly sum).
FY2025 (ended Jun 2025) was $64.6B. 2026 guidance $190B is FY2026 (Jul 2025–Jun 2026) basis.
[2] Meta initially $115–135B → revised up to $125–145B at Q1 2026 results (Apr 2026).
[3] Apple's fiscal year is Oct–Sep. 2025 actual is the value stated in the 10-K.
[4] Analyst estimate, not official Apple guidance.

Figure 1-5 Structural strengths of Japanese industry

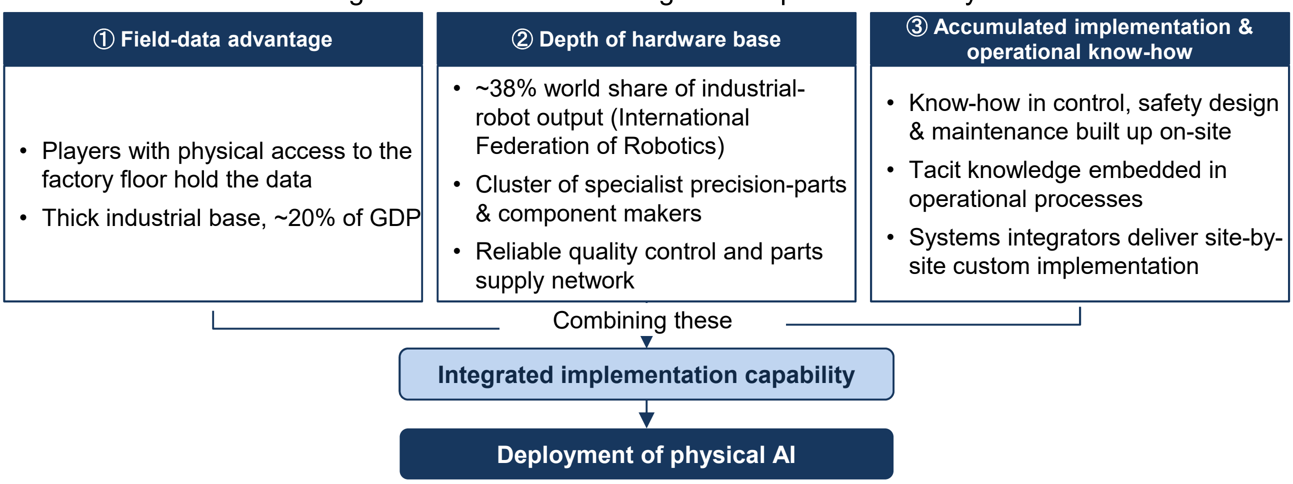

(Note) Prepared by Development Bank of Japan Inc. from various sources.

2. Physical AI Concept and Its Position in This Paper

(1) What is physical AI?

As noted in Chapter 1, in domain-specific AI, Japan's competitive edge hinges on its "integrated implementation capability" to combine field data, hardware, and operational know-how. This is critical in physical AI, where AI operates in physical space (Figure 2-1). Physical AI integrates sensing (perception), AI-based interpretation (decision), and robotic or machine execution (action). Unlike generative AI, which processes information digitally, physical AI executes tasks physically through industrial robots, humanoids, and autonomous mobile robots. While the concept remains unsettled, the Ministry of Economy, Trade and Industry (METI)'s "AI Robotics Strategy" (March 2026) frames it as a system integrating computing, control, actuation, and perception. And Gartner calls it " Physical AI is the practice of designing physical AI systems that directly interact with and operate within the physical world. These systems manipulate objects, move through space or sense physical phenomena."[2] Please note that despite differing emphases, both definitions imply that AI models alone are insufficient and require hardware integration.

[2] Gartner®, "Top Strategic Technology Trends for 2026: Physical AI", Erick Brethenoux et al., 18 October 2025. GARTNER is a trademark of Gartner, Inc. and its affiliates.

Figure 2-1 Physical AI: organizing the concept from three perspectives

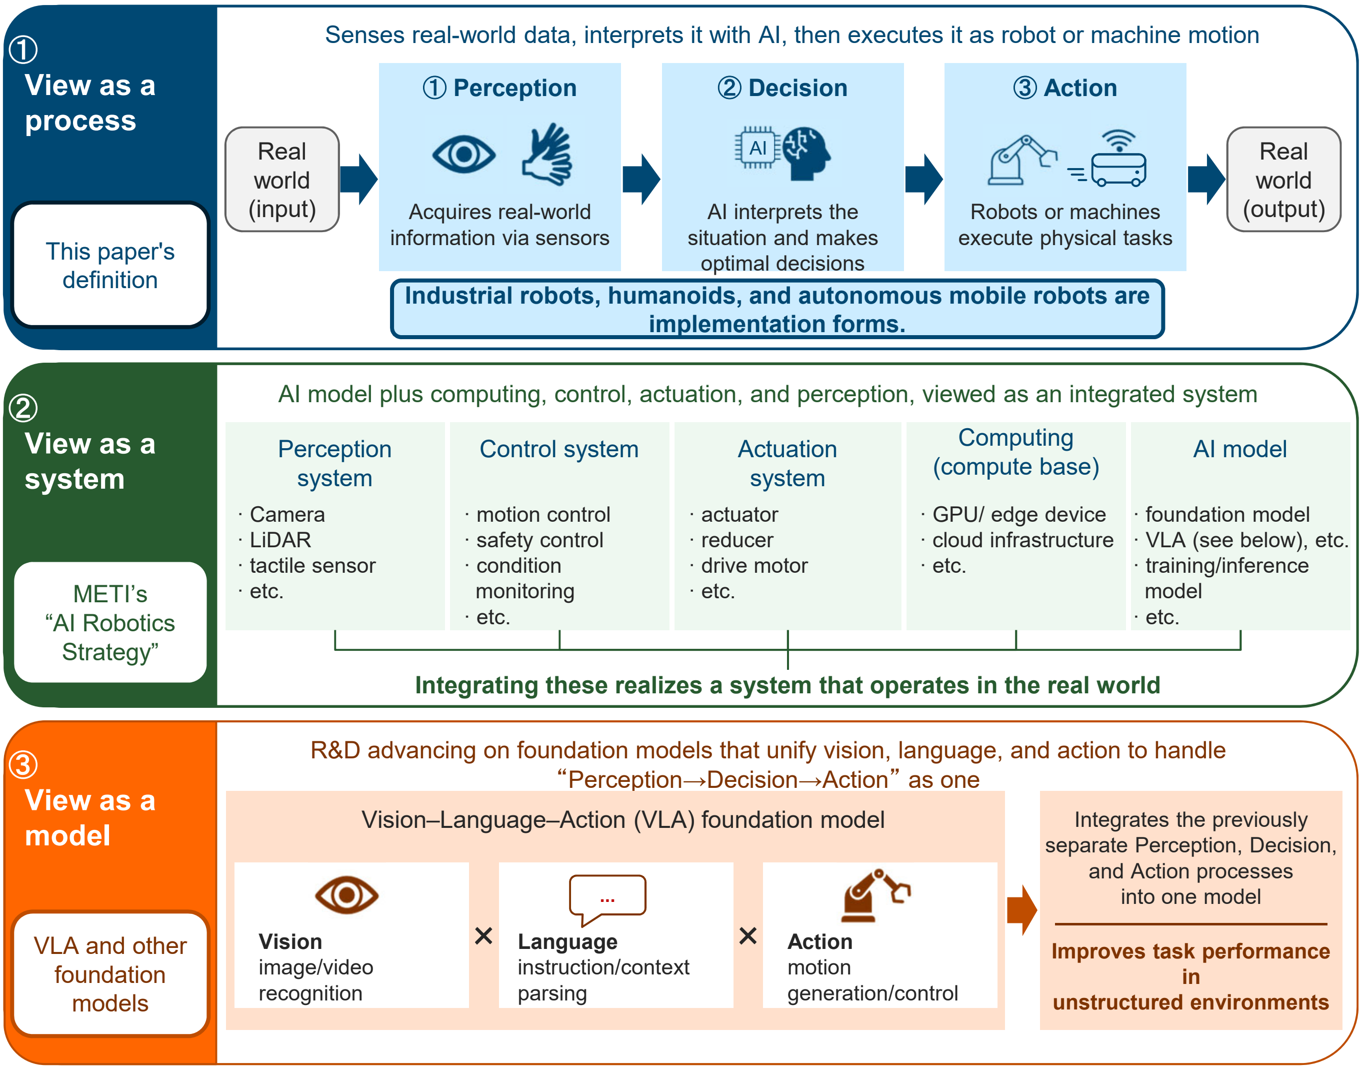

(Note) Prepared by Development Bank of Japan Inc. based on METI's "AI Robotics Strategy" (March 2026) and other sources.

(2) Positioning of This Paper

This paper treats physical AI as "a domain, centered on robotics, where integrated AI and hardware perform physical tasks on-site." Its initial market, however, is not the high-precision, high-speed manufacturing handled by existing industrial robots, but domains such as material handling, quality inspection, and labor-intensive routine work that demand situational judgment and environmental adaptation rather than precision. Such work exists widely across industries, yet conditions differ by site; the key to market formation is finding common challenges that can be bundled into a common baseline for deployment. This paper focuses not on strict definition but on the value chain structure, implementation economics, and policy issues around this domain. As a technical core of physical AI, R&D is advancing on foundation models, such as the Vision–Language–Action (VLA) model, which integrate vision, language, and action to unify perception → decision → action—a key technical foundation for raising task performance in unstructured environments.

3. Physical AI Value Chain: Structure and Bottlenecks

(1) Contrast with the IT Value Chain

Understanding social implementation of physical AI structurally is aided by contrast with the mature IT value chain (next page, Figure 3-1). In the IT value chain, stages from consulting & strategy through system design, development & deployment to operation & maintenance are differentiated, and an industry structure is established in which IT consultants, SaaS vendors, and systems integrators ("SIers") play defined roles. Recently, AI application services have been noted as a risk of substituting SIer functions, but overall the division of roles among players is clear. The platform layer—cloud infrastructure and SaaS—has been shaped mainly by U.S. hyperscalers and SaaS vendors.

(2) Structure of the Physical AI Value Chain

In contrast, the physical AI value chain has four stages, but maturity differs greatly from the IT side: (1) consulting & strategy, by robot SIers/consultants and users; (2) hardware / foundation models & apps, mainly by robot makers; (3) implementation & integration (SIer / line builder)—process design, integration, on-site tuning, and change management; and (4) business operation, where users drive continuous improvement via data, learning, and feedback. The key difference is that the third stage—robot SIer / line builder functions—is not yet established as an industry. METI's "AI Robotics Strategy" (March 2026) also makes fostering SIers a key measure, precisely where Japan's "integrated implementation capability" applies. A key bottleneck to social implementation is the absence of entities capable of managing the process from requirements definition to on-site coordination.[3]

In the U.S., in-house demand lets firms cover development through implementation via vertical integration, unlike Japan. Chapters 4 and 5 examine the SIer problem and JV-based implementation.

[3] Japan Robot System Integrator Association, "Reiwa 5 Manufacturing Base Technology Survey (Structural Issues of the Robot SIer Industry and Strengthening Its Business Base for Accelerating Robot Adoption) Report" (2024/3). METI, "On Future Robot Policy" (2025/6).

Figure 3-1 IT and Physical AI value chain comparison

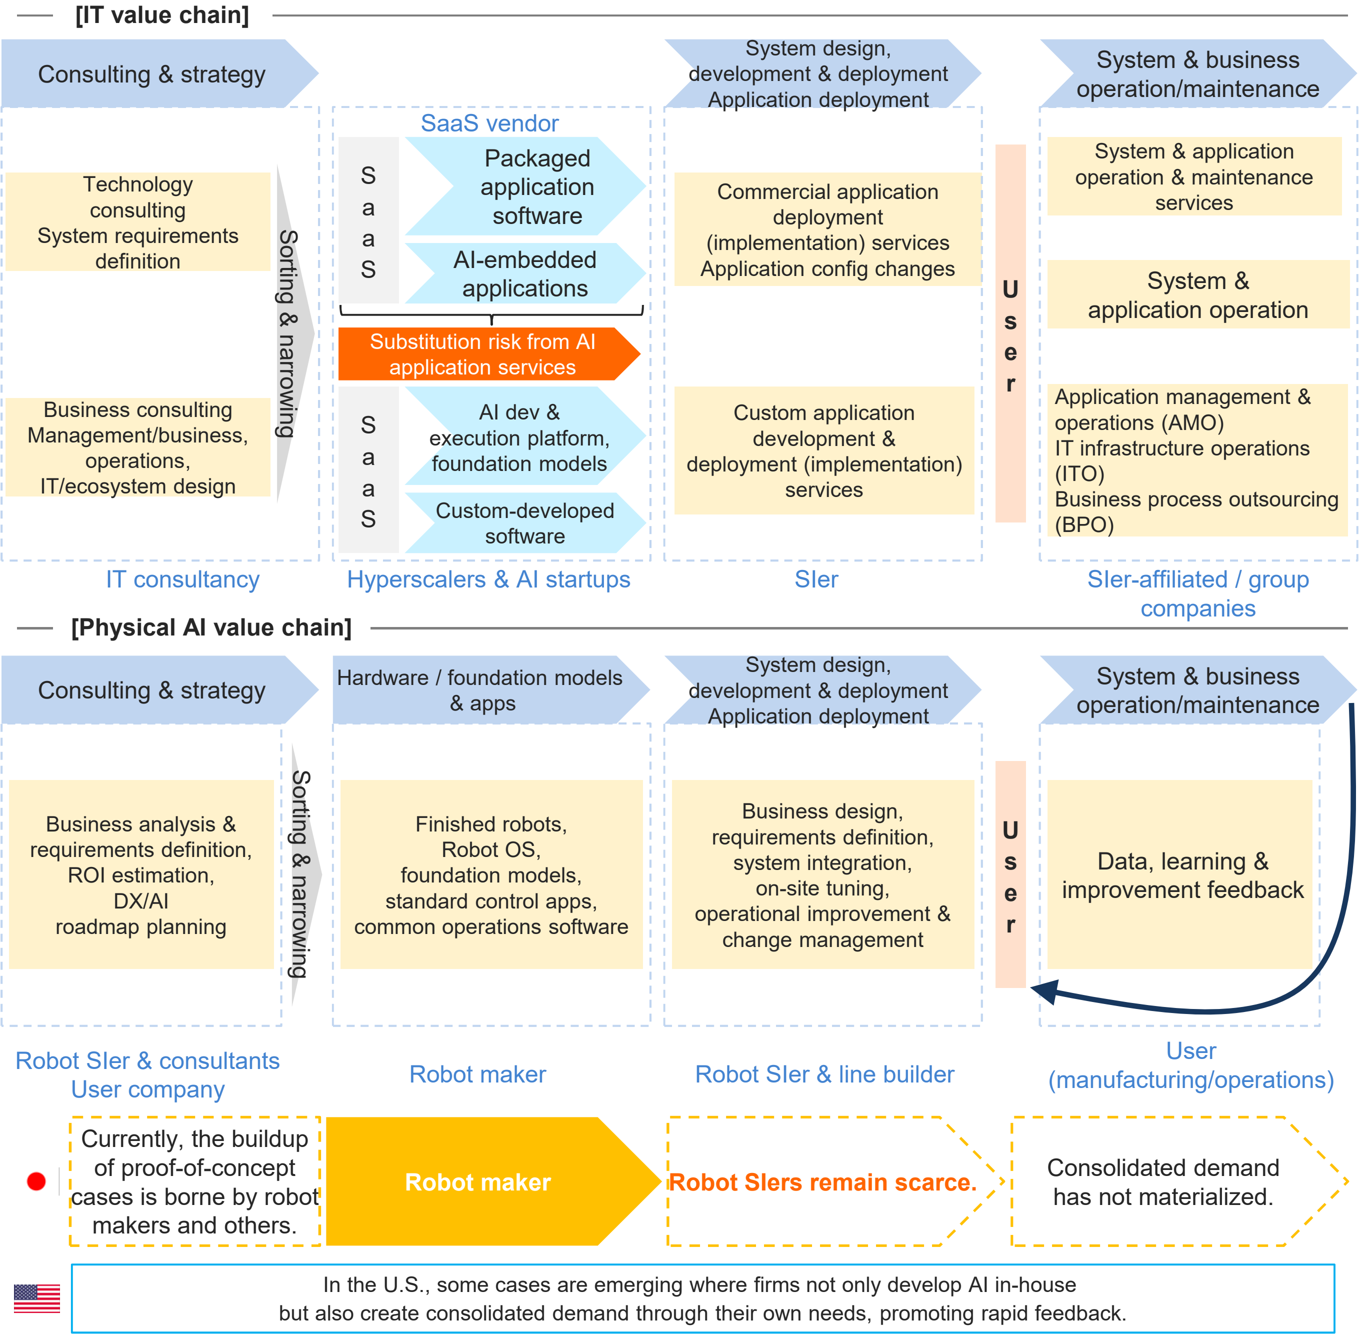

(Note) Prepared by Development Bank of Japan Inc. from various sources.

4. The Need for SIer Functions and the Absence of an SIer Industry

(1) Structure of the Problem

The physical AI value chain has a gap in implementation and integration. This chapter separates SIer functions from the SIer industry. SIer functions[4] are indispensable, yet many deployment challenges users face stem from their insufficient provision.[5] In-house provision by makers is rational, but while user deployment barriers persist, the market will not expand.

[4] The four SIer functions: requirements translation, implementation integration, on-site adaptation, and continuous learning. For definitions, see the notes to Figure 4-2 on the next page.

[5] Keidanren (Japan Business Federation), “Recommendations on Japan’s Robot (AI+) Strategy” (March 2026).

(2) SIer Industry: Structural Factors Making It Hard to Establish

The robot SIer business is man-month based, requires site-by-site customization, and is hard to scale. Demand is fragmented, maintenance revenue stays with manufacturers, and AI advances make knowledge hard to keep current (Figure 4-1). AI-driven code generation[6] also enables in-house routine development; so, future SIer demand may consolidate around specialized, on-site domains such as physical AI.

(3) SIer Functions: Current Providers

That said, SIer functions have not disappeared. Each player leverages its strengths individually and achieves certain results (Figure 4-2). The robot SIer covers functions most broadly, but scales poorly and cannot keep up with future deployment growth. Other players have distinct strengths, but none can consistently cover all functions. Fostering SIers is an important policy,[7] but in parallel it is meaningful to design a framework realizing SIer functions by combining each player's strengths. Chapter 5 discusses one such direction: implementation via a JV.

[6] So-called vibe coding: a development method in which AI automatically generates code from natural-language instructions.
[7] METI, “AI Robotics Strategy.” (March 2026)

Figure 4-1 Ease of attracting growth investment by business model

	SaaS	IT SIer	Robot SIer
Revenue model	Subscription (monthly)	Person-month contracting	Person-month + on-site staffing
Scalability	High (low marginal cost)	Low (staffing needed per project)	Even lower (custom work per site)
Revenue predictability	High (manageable via churn rate)	Moderate (backlog gives some visibility)	Low (demand fragmented by site)
Knowledge reuse across cases	Easy (same product to many)	Somewhat (reuse of industry knowledge)	Difficult (site conditions differ each time)
Recurring revenue after deployment	Yes (monthly billing continues)	Limited (maintenance contracted separately)	Little (maintenance tends to stay with makers)
Public-market valuation	30+ listed on TSE	20+ listed on TSE	Zero pure-play listings

(Notes) Prepared by Development Bank of Japan Inc. from various sources. SaaS = Software as a Service. Robot SIer = a firm specializing in the design, deployment, and tuning of robot systems. Of Japan Robot System Integrator Association full membership, 217 firms are unlisted SMEs.

Figure 4-2 SIer functions: current providers

	① Requiremant Translation	②Implintegration	③ Field adaptation	④ Continuous learning	Strengths & limits
Robot SIer	◎	◎	◎	○	Covers all four functions most broadly, but scales poorly and lags growing deployment demand
Robot maker	○	◎	◎	○	Strong at integrating and adapting its own products, but other vendors’ products and AI fall outside its scope
IT SIer	◎	○	○	◎	Strong in system design and data infrastructure; hardware and on-site work lie outside its expertise
Consulting firm	◎				Turns management issues into technical requirements, but rarely engaged beyond that stage
AI vendor	○	○	○	◎	Model development and continuous learning are core strengths, but on-site hardware and process knowledge are limited

(Notes) Prepared by Development Bank of Japan Inc. from various sources. SIer functions are defined as follows: ① Requirement translation: converting business issues into technical requirements; ② Implementation integration: unified deployment of hardware, software, and field; ③ Field adaptation: tuning and safety handling; and ④ Continuous learning: model improvement from operational data. The ◎ and ○ ratings show general tendencies; the scope handled varies by firm and by project.

5. JV to Realize SIer Functions

(1) Options for User Companies

As outlined in Chapter 4, SIer functions are essential, but the industry providing them is immature and each player's strengths are dispersed. From a user's view, even if physical AI implementation is outsourced to an IT vendor, missing SIer functions tend to stall it at the proof of concept stage. Yet few firms can invest at large scale alone, and externally led reform via acquisition/investment can undermine management autonomy. Robot adoption also incurs initial investment plus subscription costs, like AI model updates and maintenance; so, total-cost design across cost-sharing options is key. A JV, rather than ceding control to one side as in outsourcing or acquisition/investment, lets a technology-holding firm and a domain-holding user implement jointly (Figure 5-1). It sets investment burden, data ownership, and responsibility by contract, defines data access rights in a shareholders' agreement, and speeds decisions by avoiding inter-department coordination.

JV partners could include an AI vendor, IT SIer, robot maker and others; here we frame it as AI vendor × domain user (e.g., manufacturing). Robot selection must be assessed through cost–benefit analysis of the user's field needs, and fit with the AI vendor, track record, and adoption/maintenance costs; therefore, the base scheme does not include robot makers as JV partners (next page, Figure 5-2). Once robot selection is clear, including a robot maker as an equity partner may be beneficial. In operation, AI vendor engineers pair with the user's field staff and the JV takes the core SIer functions role; hardware installation is supplemented via alliances with robot makers.

Figure5-1 Comparison of business forms in Physical AI implementation

	Outsourcing (contracting)	Acquisition/investment (M&A)	Joint venture (JV)
Overview	• User company outsources to an AI vendor/SIer	• User company acquires an AI/robotics firm	• User company and AI vendor co-invest
Capital relationship	• None	• One party acquires shares	• Both parties invest
Data access & ownership	• Scope set by each contract	• Data held by the acquired firm	• Both contribute; accumulated in the JV; ownership set by shareholders' agreement
Decision-making structure	• Client orders, contractor executes	• Follows the acquirer's management policy	• Joint decisions among investors • Coordination costs tend to arise
Vendor lock-in	• Tends to depend on the contractor's tech/products (mitigable via split orders)	• Depends on the acquired firm's existing tech (complementable via external alliances)	• Design can assume alliances with multiple firms • Partner selection needs care (dissolution/change is not easy)
IT/AI vs. robot coverage	• Depends on contract form: prime-contractor bundle or individual orders	• Centered on the acquired firm's coverage • Gaps filled via external alliances	• JV handles IT/AI in-house; robots covered via business alliances
Continuity of model improvement	• Depends on contract term; ends if not renewed	• Continues as the acquired firm's business; resource allocation may shift with integration	• JV business; risk of stalling if investors' policies diverge

(Note) Prepared by Development Bank of Japan Inc. from various sources.

Figure 5-2 Physical AI implementation: image of the JV structure

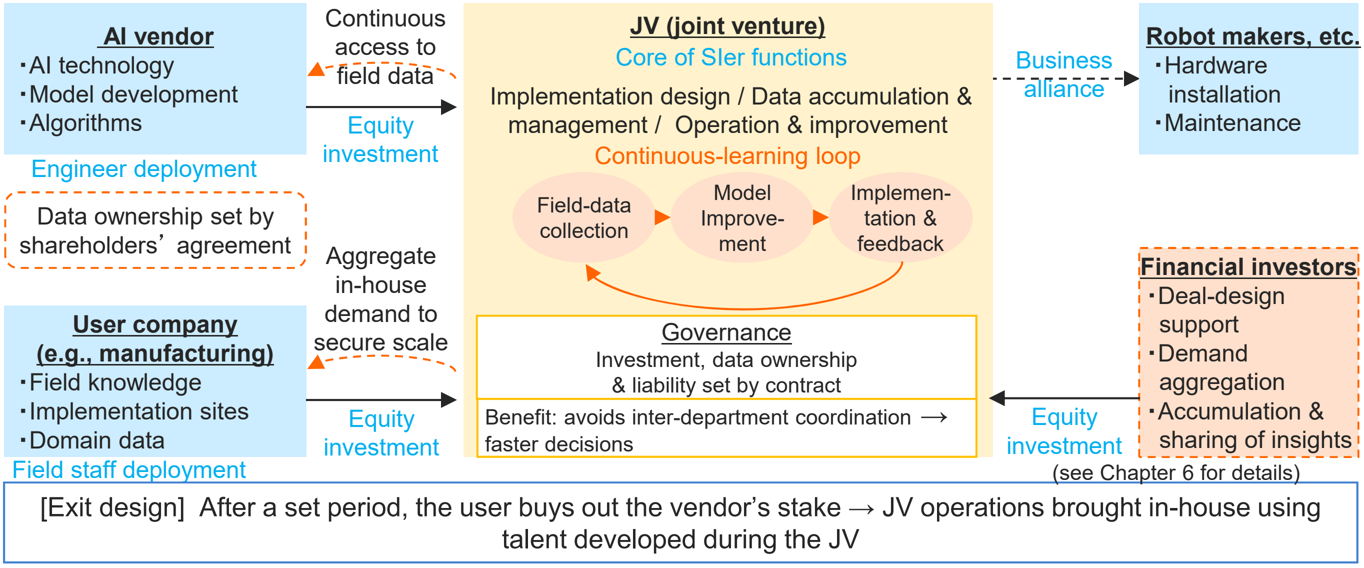

(Note) Prepared by the Development Bank of Japan Inc. from various sources.

Figure 5-3 JV lifecycle & exit design

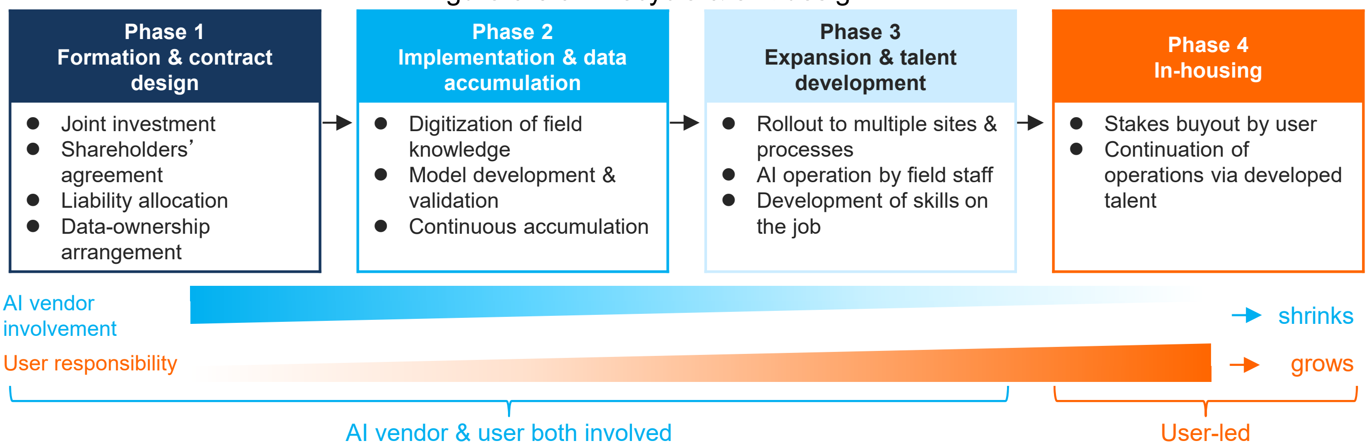

(Notes) Prepared by Development Bank of Japan Inc. from various sources. Each phase's duration varies by project. Phase 4 may not be reached; some cases continue.

(2) A Mechanism to Turn Field Knowledge into Data

At many manufacturing sites, procedures and know-how exist but go unrecorded as data. With outsourcing, data sharing stays project-based and stops when the contract ends. A JV builds continuous data accumulation and usage into the organization and fosters the daily collaboration and skill transfer that outsourcing cannot. It systematically makes Japanese manufacturers' accumulated field knowledge AI-usable. By consolidating sites and processes, the user eases demand dispersion, while the AI vendor keeps accessing, as co-owner, field data arising only through the JV—benefits both sides could not otherwise gain. As an exit design, after a set period the user could buy out the vendor's stake and run operations with talent developed during the JV (Figure 5-3).

On data ownership, the field knowledge captured by physical AI stays with the user company, not the AI vendor or robot SIer. Vendors and robot makers instead differentiate through general-purpose operational capabilities distinct from any one company's field knowledge (e.g., tidiness, handling foreign-object contamination), learning efficiency, and cost. Only when such user-specific knowledge and a general-purpose technology base combine does international competitiveness emerge, opening the way to expand a Japan-born physical AI model overseas. The JV secures this integration.

(3) Challenges Toward Realization

SIer and user-company JVs are already emerging. And overseas, tie-ups with AI vendors are appearing (Figure 5-4). In Japan, such AI vendor/user combinations remain few, with challenges: separating vendor IP from JV-accumulated data, revising functional scope as AI improves, and coordination with existing SIers and vendors. The basic approach is to validate case by case and extend successes to other industries and regions, though rapid spread within an industry is possible amid worsening labor shortages. Institutional support for such JVs is covered in the next chapter.

6. The Role of Policy and Finance for Social Implementation of Physical AI

(1) Absence of Demand-Side Policy

The Ministry of Economy, Trade and Industry's "AI Robotics Strategy" (March 2026) targets supply-side capacity but offers no measures to reduce users' upfront investment. The U.S. is similar. China backs both supply and demand via local subsidies, tax incentives, and low-interest loans, creating a policy gap (Figure 6-1). In Japan's GX field, a scheme subsidizing up to 50% of user capital investment, conditional on 100% decarbonized electricity,[8] already operates.

(2) Designing a Demand-Side Subsidy

A similar demand-side subsidy could suit physical AI, conditioned on domestic control of training/inference data (data sovereignty). Applied regardless of product nationality and grounded in economic security to keep field data from flowing abroad, it can be designed not to conflict with the WTO Subsidies Agreement's local-content ban. The proposed JV-type model keeps data domestic and avoids single-vendor dependence. If such subsidies spur JV formation, a virtuous cycle follows: supply-side support → JV formation → faster deployment (Figure 6-2).

[8] The GX Strategic Zone scheme: see Ministry of Economy, Trade and Industry, "GX Strategic Zone Scheme" overview page, Decarbonized-Power Regional-Contribution Investment Promotion Project. (https://www.meti.go.jp/policy/energy_environment/global_warming/gx_strategy_area.html)

Figure 5-4 Examples of JVs related to SIer function

Type	Partners	JV name	Timing	Overview
IT SIer × Manufacturing	SCSK × NGK Spark Plug (Niterra)	SCSK Niterra IT Solutions	2025	Manufacturer's IT operation/maintenance carved out to a JV, fused with SCSK's IT operations capability
AI Tech × Shipbuilding / Shipping	JDSC × Tsuneishi Shipbuilding × Mitsui & Co.	seawise	2022 founded	Built a ship-data platform. After deployment on over 200 ships, as the business became self-sustaining, JDSC transferred its stake (2025)
AI Robotics × Manufacturing	Intrinsic (Alphabet) × Foxconn	(Undisclosed)	2025	AI robotics software × electronics manufacturing. Jointly developing AI-driven robotics at a U.S. plant
AI Vendor × PE Alliance	Anthropic × Blackstone, etc.	(Undisclosed)	2026	AI implementation-services firm established. Supports AI adoption at client companies via an embedded-engineer model

(Note) Prepared by the Development Bank of Japan Inc. from various sources.

(3) Building the implementation foundation

Developing supply-side technology and demand-side capacity in parallel, realizing the SIer function via frameworks such as JV, and building the institutional foundation is the next policy task. Bundling common issues by industry or region eases adoption even when single-company costs are heavy. Accumulating sector- and region-level patterns for horizontal rollout prepares for an adoption surge from labor shortages, as normal-time verification enables faster crisis deployment. Recording Japanese manufacturing's on-site expertise as AI-usable data is the starting point for physical AI's social implementation.

Figure 6-1 Physical AI / robotics policy: international comparison

	Japan	US	China
National strategy	"AI Robotics Strategy" (March 2026) Robot makers, foundation models, domestic components	No robotics-specific national strategy **yet established** (as of July 2025)	15th Five-Year Plan (2026): robotics as a strategic emerging industry, embodied AI as a future industry
Supply side: R&D and industry support	- FY2026 physical AI-related: 387.3 billion yen - New Industry and Industrial Technology Development Organization (NEDO): robot software platform (3 yrs, 10.3 billion yen)	- CHIPS and Science Act (2022): $52.7 billion for semiconductor mfg and R&D - NSF: robotics research, $30–50 million/yr (private VC active in humanoids, e.g., Figure AI)	- National VC fund: 1 trillion yuan (20 yrs) - Local govts: industry funds, R&D/capex subsidies, low-interest loans (Beijing, Shenzhen, etc.)
Demand side: subsidies for user robot adoption	**Limited at present** Implementation roadmap (17 fields) and adoption targets green transformation (GX) field: capex subsidies exist[1]	**None at present** (as of July 2025)	Central & local demand-side support - Robot-purchase subsidy (avg. 17.5%) - Interest subsidy for equipment renewal - Digitalization tax credit (10%)

[1] GX Strategic Zone: capex using 100% decarbonized electricity (direct current [DC] & manufacturing equipment) gets up to 50% subsidy (see Note 8, previous page)

(Note) Prepared by Development Bank of Japan Inc. from various sources.

Figure 6-2 Physical AI: Illustration of user subsidies

(Reference) GX Strategic Zone user subsidies

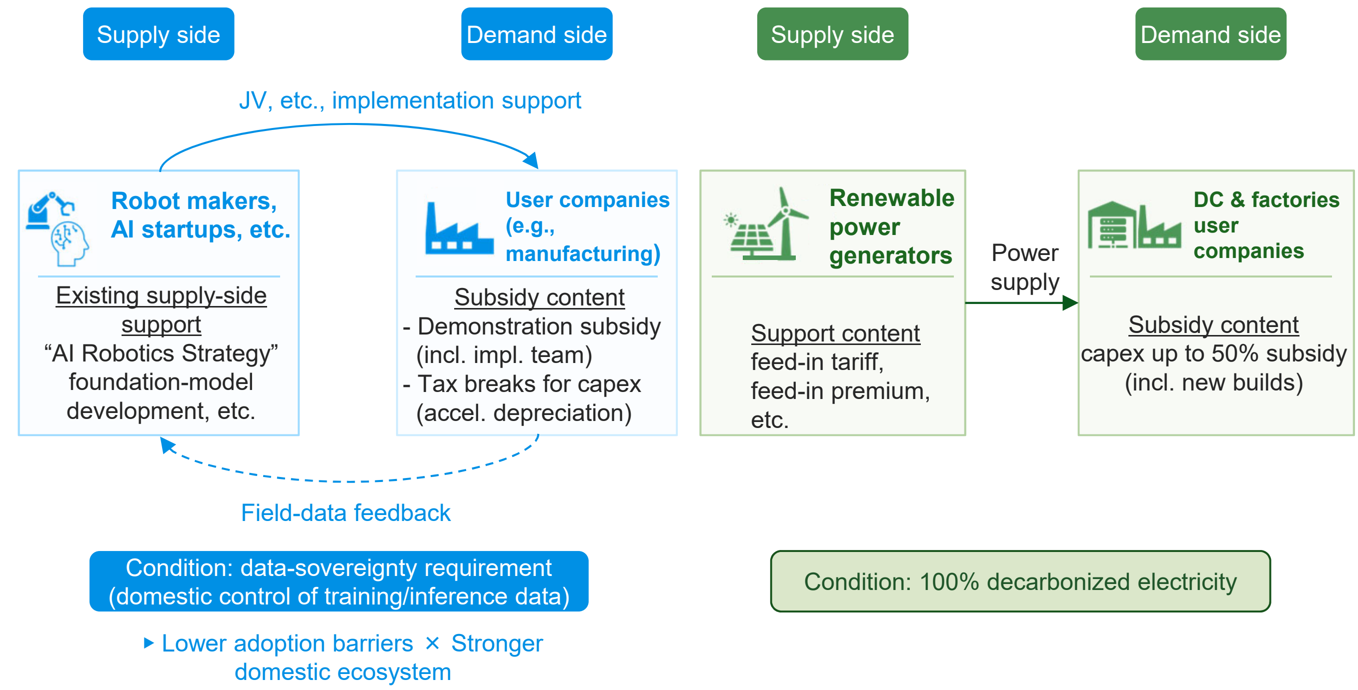

(Note) Prepared by Development Bank of Japan Inc. from various sources.

(4) Role of Financial Institutions Supporting Deployment

Physical AI's social implementation is not only a technology issue but also a business-design issue. Chapter 5 discussed the JV, in which a technology holder and a domain-holding user jointly deploy physical AI. Forming such a JV requires functions the parties alone often struggle to provide. Financial institutions, connected to many industries and companies and independent of specific vendors, are well placed to fill this role (Figure 6-3).

Function 1: deal-design support at the JV conception stage. A JV's success depends on pre-formation design: who partners with whom, in which domain, and with what roles. User companies struggle to assess AI technology, while AI vendors may not grasp users' priorities. A neutral party that structures the deal and matches partners can smooth the JV launch.

Function 2: demand aggregation. As noted in section (3), the cost burden of adopting physical AI is heavy for a single firm. Bundling common issues across industries and regions into a JV makes SIer functions easier to commercialize. Financial institutions can aggregate dispersed small and medium enterprise demand by combining wide-area deal design with local contact points.

Function 3: accumulating and sharing insights on implementation economics. The biggest adoption barrier is foreseeing return on investment. Insights on viable industries and processes, break-even scale, and fair risk-sharing contracts emerge only through cases. By systematizing and sharing them, financial institutions can provide a public good that accelerates later adoption.

The functions listed above precede investment and lending, but can unlock diverse funding, including private capital. Demand-side subsidies promote JV formation, financial institutions support deal design and demand aggregation, and users provide data and field knowledge. When these elements align, a virtuous cycle links supply-side support to faster deployment.

Starting from the shift in the generative AI market's competitive axis, this paper analyzed the shortage of SIer functions as a structural obstacle to physical AI's social implementation and proposed a JV-based deployment system with demand-side subsidies. Supply-side support strengthens technology development, but delivering technology to the field requires a mechanism integrating industry, policy, and finance. This is the next challenge for social implementation of physical AI.

Social implementation of physical AI has only just begun, and many issues remain beyond this paper's scope. The bottlenecks lie largely beyond technology, and we will continue mapping the path forward.

Figure 6-3 Role of financial institutions in physical AI deployment

Role	Details	Target
① Deal-design support	Structuring the business, matching candidate partners, designing role allocation	JV conception stage
② Demand aggregation	Aggregating common issues across industries/regions to secure profitability of SIer functions	Small and medium manufacturers, regional company clusters
③ Accumulating and disseminating insights	Systematizing insights on the economics of implementation (ROI, break-even scale, contract design)	The whole industry

(Note) Prepared by Development Bank of Japan Inc. from various sources.

Contact: Industrial & Regional Research Department, Development Bank of Japan
Tel: 03-3244-1840
e-mail: report@dbj.jp