



FROM RELIABILITY TO AI-READY: JAPAN'S ROADMAP TO MEASURABLE ENTERPRISE VALUE

Japan can accelerate into the AI era without sacrificing reliability by making three deliberate moves: refresh the core where risk and cost accumulate, renew foundations for cloud agility and data flow, and reimagine work and operations with AI-led innovation.

INSIGHTS



Japan needs to refresh aging core systems and reduce knowledge risk as experienced specialists retire.



Many IT landscapes were built for stability, not cloud agility and AI integration. That raises maintenance costs and operational risks.



Organizations should prove AI's value through one or two use cases and then scale.



Strong data foundations, integration, engineering, and operations discipline make scaling practical.



Reimagine how humans and machines work together through a product-centric operating model, supported by clear governance.

Protect what works, refresh the constraints that slow change, renew cloud and data foundations, and then reimagine workflows with AI. The Japanese industry is known for reliability and process excellence. This operational rigor has led to groundbreaking achievements in manufacturing and process engineering, widely adopted by leading companies globally.

Although reliability is a strategic asset, the AI era is changing the dynamics of competition. Speed, data-driven decision-making, and [engineering productivity](#) are now core capabilities to succeed in the emerging operating model.

In Japan, the barrier to scaling AI is not a lack of intent. Companies such as [SoftBank Robotics and Fujifilm](#) are moving quickly to apply AI in manufacturing and healthcare, and similar momentum is emerging in smart cities and financial services.

The constraint is execution: many organizations must deliver AI at scale across complex and often siloed application landscapes, while the pool of specialists with the know-how to renew them continues to shrink.

This renewal is urgent and the window to act is now short. The Ministry of Economy, Trade, and Industry first warned eight years ago in 2018 about the “2025 Cliff.” The report stated that unless organizations renew IT systems and fully embrace digital, Japan will experience a [shortfall of ¥12 trillion \(\\$77 billion\) per year after 2025](#).

In all of this, the goal is not to replace what works. Japan can keep reliability intact while restoring speed and adaptability through organizational reinvention. The path forward is structured and low-risk: refresh what must change, renew foundations for agility, and reimagine work with AI while keeping reliability nonnegotiable.



The reliability gap: When stability slows outcomes

It's not as if Japan, like Europe and the US, has not made inroads into foundations that are ready for AI. Many Japanese organizations, such as Fujitsu, have updated their technology stacks for a [cloud- and data-driven era](#), transitioning to microservices architectures that support modern workloads.

But for many, much of their IT landscape remains heavily customized and connected to disparate systems, both old and new. This “spaghetti” of interfaces and processes does not readily work with AI technologies. Because the business just worked in the past and IT systems kept running, many organizations focused on refreshing the front end through web-based applications — creating a stronger presentation layer but without the foundational speed needed. As a result, digital touchpoints improved, yet cycle times did not, because constraints remained in the surrounding systems and process dependencies.

Lack of AI integration is one challenge. Keeping the lights on is also getting more difficult. Most Japanese universities now teach Python and newer application-focused languages; few graduates want to learn older languages such as COBOL or Pascal. As experienced specialists retire, they take critical business and technical

knowledge with them. This increases operational risk and slows the pace of change, making outcome-focused renewal programs urgent.

There are budgetary concerns too. Government and industry analyses note that many enterprise systems in Japan are decades old and that, for a significant share of firms, maintenance consumes the bulk of IT spend. As Nandan Nilekani, Chairman of Infosys, said in a [recent keynote](#), organizations want to invest far more in innovation and new capabilities such as AI, but that shift is hard without first simplifying what they already run on today. For example, a Japanese bearing manufacturer highlighted challenges in its e-commerce site, including system aging, rising maintenance costs, and the need to update the operating system and middleware. Moving to cloud was seen as nonnegotiable to achieve readiness for further enterprise transformation.

Older IT environments can also elevate disruption risk. Incidents and breaches can go undetected for months, increasing financial and reputational exposure.

Further, IT in Japan has historically been treated more as a risk-control function than as an invention engine. That shapes governance, ownership, and pace. Decision authority often sits with business leaders, with procurement and risk functions playing a strong role in vendor selection. CIOs and CTOs therefore need a clear value case — one that proves reliability will not be compromised and that benefits will be visible early, not only after multiyear programs.




An outcome-focused AI roadmap for Japan

An AI-first organization delivers business value faster, increases productivity, and drives efficiencies at scale.

To help organizations in their AI journey, Infosys has developed an AI-first value framework, a comprehensive suite of services to guide and implement the AI transition (Figure 1).

Figure 1. Infosys AI-first value framework comprises six key pillars



Source: Infosys

The framework aligns leaders on outcomes first and then translates those outcomes into the foundations and execution needed to deliver them. Outcomes may include improving customer experience through more personalized products, reducing cycle times in critical processes, or lifting employee productivity by simplifying work. In practice,

the journey works best through three moves: refresh the constraints that lock up budgets and knowledge, renew foundations for cloud agility through data and integration, and reimagine processes and experiences with AI-led innovation, with responsible AI, security, and compliance built in.

Japanese enterprises can use the framework to align on outcomes and move quickly toward:

	Productivity and speed through AI-assisted engineering, modern delivery discipline, and shorter change cycles.
	Experience and growth by decoupling journeys from constraints, using standard interfaces, and ensuring data availability.
	Data and integration that make AI usable in real operations, with accessible, live data and low-latency interfaces.
	Resilience and reliability through observability and continuity patterns that protect critical operations as change accelerates.
	Risk, security, and compliance by design , with governance and controls that build trust and enable faster adoption.
	Workforce readiness by capturing knowledge, reducing dependency on scarce expertise, and upskilling teams for AI-assisted ways of working.

Productivity and speed are often where leaders can prove value quickly. AI-assisted engineering can compress delivery cycles, reduce rework, and lower dependence on scarce skills, especially when combined with strong delivery discipline.

Generative AI and more advanced autonomous agents can accelerate tasks such as [code generation, refactoring, testing, and documentation](#). Proof-of-concept timelines shrink when requirements documents can be converted into working code faster, and when test and release cycles are streamlined. AI-assisted engineering also supports knowledge capture and helps teams understand complex landscapes.

Renewing foundations for cloud agility requires stronger data flow and integration that unlock end-to-end outcomes. This means ensuring that data is accessible, readable by AI systems, live, and delivered with low latency to systems doing the inference. AI systems should integrate with older software components through application programming interfaces (APIs). As organizations onboard new models and infrastructure, these interfaces preserve flexibility and reduce vendor lock-in.

AI safety and security are also paramount. One Japanese gaming enterprise worked with Infosys to create a content moderation system that utilized large language models to determine whether content met business policy requirements. Key was ensuring all outputs, filtered for hate speech, personal privacy, and copyright, supported downstream automation by enabling rule-based actions, human review prioritization, and auditable transparency for compliance.





A quick maturity check to decide where to start

Before launching proofs of value, leadership teams can run a simple assessment across five areas: data foundations, integration and APIs, engineering and operations discipline, governance and trust, and workforce readiness.

This quick check also clarifies what to refresh first (core constraints and knowledge risk), what to renew (cloud, data, integration), and where to reimagine workflows with AI.

Scoring each area from 1 to 5 clarifies where AI will stall today and which upgrades will unlock the biggest outcomes first (Figure 2). It also helps align business, IT, risk, and procurement on what will be proven in 60 to 90 days and what must be built to scale.

Figure 2. Maturity rubric for AI readiness

Area/ Maturity	1 (Ad hoc)	2 (Emerging)	3 (Defined)	4 (Scaled)	5 (Optimized)
Data foundation	Siloed and inconsistent	Partial accessibility	Governed with metadata	Real-time, enterprise-wide access	Fully integrated, feedback-driven
Integration and APIs	Point-to-point integrations	Basic APIs, early integration efforts	API-first and reusable services	Event-driven architecture, scalable integration	Agent-enabled orchestration across systems
Engineering and operations	Manual experimentation	Early pipelines, limited standardization	Standard MLOps, CI/CD for models	Automated pipelines and model life cycle management	Self-optimizing pipelines and continuous learning systems
Governance and trust	No formal policies or controls	Initial guidelines (risk, privacy)	Defined policies and compliance processes	Continuous monitoring, auditability, risk controls	Adaptive, real-time governance
Workforce readiness	Limited awareness, no formal skills	Small trained teams, pockets of expertise	Broader training, defined AI roles	Widespread AI fluency, cross-functional adoption	AI-native workforce, humans + agents teams

Source: Infosys



Japan's prove-then-scale playbook

A helpful lesson comes from a retailer that moved its website to [SAP Commerce Cloud](#) and improved the front-end experience. The project highlighted the importance of refreshing core back-end constraints in parallel with the user interface to unlock the full potential of end-to-end improvements.

To move ahead, organizations should start with one or two outcomes that leaders care about, and make value improvements visible to executives in proofs of concept before scaling more widely. Value metrics here include cycle-time reduction, productivity per user, and service reliability.

The playbook we advocate for Japanese organizations is composed of three stages:

STAGE 1

Reimagine and prove value in 60 to 90 days

- Pick one or two outcomes and define metrics that leaders care about.
- Bring business owners, risk, and procurement into the proof early, and agree what success looks like.
- Run a contained proof of concept that is safe, measurable, and visible. In Japan, proof builds permission to move ahead.

STAGE 2

Renew foundations to scale

- Create data access and quality for the chosen domains.
- Build standard interfaces and integration discipline.
- Ensure security and compliance by design.

STAGE 3

Refresh the core and make delivery repeatable

- Move toward [product-centric ownership](#) for priority domains.
- Set up platforms so that disparate teams don't reinvent everything.
- Provide high-level governance and responsible AI controls that speed decisions rather than slow them.

For example, outcomes might be improving release frequency by more than 30%, while reducing lead time by 40%, the average time to resolve incidents by 25%, and cost-to-serve by 10%.

As organizations move from one of two use cases to scale, ensuring data readiness is key. Organizations will have to put in place [master data management and control, and metadata catalogs for effective data governance](#). Doing this, along with deploying the right integration architecture, ensures vital domain knowledge isn't lost, and that it can be processed quickly enough for time-sensitive operational processes. For example, one Japanese sports arena's journey to green energy required a centralized data platform to integrate, standardize, and visualize sustainability-related data so the company could track carbon footprint performance, and drive data-driven decarbonization initiatives. By centralizing data and enabling automated data pipelines to process large volumes of operational data in near real time, the solution empowered business teams to meet rigorous efficiency goals.

Making delivery repeatable is what helps you move from pilots to results at scale, so stage 3 needs real focus.

Product-centricity helps organizations track the flow of value, creates alignment between teams for priority domains, and encourages engagement around measurable goals. Instead of siloed business units shaping systems independently, creating product-focused teams with cross-functional support ensures the right system gets built. Progress can then be measured by an effective set of [objectives and key results](#). Key will be to also put in place data product owners, [an appropriate funding model](#), and the right evaluation and monitoring cadence.

A well-tested engineering platform means the new architecture stays adaptable as foundations scale, while ensuring critical systems remain operational and downtime risk stays low.

The final step of stage 3 is key. Case-by-case governance tools should be replaced by a centralized governance layer. [IDC FutureScape 2026](#) expects that by 2027, 40% of global enterprises will replace fragmented oversight with unified AI governance frameworks to manage risk, enhance trust, and enable innovation as more autonomous AI is adopted. To help, organizations can align a minimal responsible AI controls checklist across privacy, logging, evaluation, monitoring, and auditability (Figure 3). According to our paper, [A new approach to safe agentic AI](#), good governance requires built-in, continuous controls across the entire life cycle, not one-off checks.

Figure 3. Responsible AI controls checklist

Control	Checklist
 Privacy	• Strict data access boundaries • Data minimization and masking • Data leakage containment
 Logging	• End-to-end agent telemetry • Immutable logs for all decisions and interactions • Traceability across multi-agent workflows
 Evaluation	• Predeployment testing and red-teaming • Evaluation across prompts, models, systems, and outcomes • Continuous accuracy, safety, and business alignment assessments
 Monitoring	• Continuous observability and performance monitoring • Drift and anomaly detection • Feedback loop integration for ongoing improvement
 Auditability	• Full life cycle traceability • Audit-ready evidence for compliance and governance • Controls embedded into a centralized governance layer

Source: Infosys

The goal then is to be AI-first, much like FANUC and SoftBank Robotics in Japanese manufacturing, or Mizuho Financial Group, who use a [center of excellence or “AI promotion office”](#) to drive business strategy and get the whole organization behind the AI transformation.



How Japan can move faster without breaking reliability

Global competitiveness is being redefined by who can scale AI safely and quickly.

However, a [Reuters report](#) found in 2024 that many Japanese executives have not yet seen the benefits of AI-ready speed and execution. A quarter have already integrated AI into their business, but as many as 40% are still on the fence.

To move ahead boldly into 2027, Japanese leaders should create a workforce ready for AI, as this underpins refresh, renew, and reimagine. The [AI Business Value Radar 2025 report](#) found that preparing employees to understand AI and engage in changes can improve success rates by 18 percentage points. This also means balancing the Japanese philosophy of [no tolerance for failure](#) with a mandate for fast, cyclical experimentation that scales.

Japan will do well to ensure reliability stays nonnegotiable, using this core cultural strength to its advantage. This should be paired with the growing realization that good AI outcomes require foundations that are ready for AI and the discipline to execute. For Japan specifically, the prove-then-scale playbook can increase the CIO's standing in delicate power dynamics where business units tend to have more clout.

Of course, innovation is not needed everywhere. Some IT systems are still adequate for certain processes and should be protected. The focus is to refresh what constrains progress, renew foundations for cloud agility, and reimagine work with AI where it creates measurable value. That combination prepares Japanese businesses for the AI era while keeping reliability nonnegotiable.



Authors

Ramesh Narayan, Infosys Knowledge Institute, Bengaluru

Harry Keir Hughes, Infosys Knowledge Institute, London

Contributors

Sachin V. Kulkarni, Infosys Japan, Tokyo

Arumugam Vallinayagam, Infosys Japan, Tokyo

Jo Ando, Infosys Marketing, Tokyo

Shreedeeep Panicker, Infosys Marketing, Hyderabad

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at infosys.com/IKI or email us at iki@infosys.com.

For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.