REPORT

AI sovereignty

A strategic imperative for European industry

ALEPH ALPHA

Roland Berger

Management summary

Driven by advances in artificial intelligence (AI), industrial operations are evolving – from product development to execution on the factory floor. While global competitors such as the US and China accelerate their investments and innovation, European industry remains deeply reliant on foreign technology and infrastructure. More than 80 % of digital products, services and intellectual property (IP) used in the EU are provided by non-European companies. With only three major AI models developed in Europe in 2024, compared with 40 in the US, Europe risks falling behind in a technology increasingly central to industrial value creation.

At the core of this challenge lies the issue of sovereignty. As industrial processes become more autonomous, data-driven and integrated, companies must maintain full control over the AI systems powering these operations. This means building trustworthy, explainable and domain-specific AI that complies with European regulatory and operational standards, as well as creating secure, scalable AI infrastructure suitable for local, hybrid or air-gapped environments.

Roland Berger and Aleph Alpha are actively supporting leading industrial players in navigating the AI transition. Drawing on hands-on implementation experience and close collaboration with policymakers and pioneers across sectors, we outline how European companies can move beyond pilot projects toward impactful, scalable deployment. The most successful firms clearly define their control and compliance requirements, prioritize high-value use cases and integrate modular, sovereign AI into their core engineering and production systems. In this context, AI sovereignty is not just a strategic advantage – it is a prerequisite for sustained industrial leadership.

Contents

1 A new reality for EU industry

Why AI sovereignty is the solution

Europe's economic backbone is facing a fundamental shift in the competitive environment, driven by two converging forces. On the one hand, increasing geopolitical uncertainty and tech nationalism are putting pressure on global trade and fragmenting established value chains. Recent US tariffs on European steel and aluminum exports, for example, are already leading to production relocations to Asia and North America – to the detriment of EU industry. Meanwhile, Europe's broader industrial sector is becoming increasingly dependent on foreign technology providers for core digital services and infrastructure. According to the Draghi report on European competitiveness, the EU relies on foreign countries for more than 80% of digital products, services, infrastructure and IP.

At the same time, European companies are under increasing performance pressure as the pace of innovation accelerates. Global competitors are drawing on advanced digital capabilities – such as autonomous systems, predictive manufacturing and generative design – at growing speed and scale. The US holds a dominant position in AI, outpacing even China; according to data from Stanford University, US institutions developed 40 major AI models in 2024, compared with 15 in China and only three in Europe. China's comparatively modest investment figures, despite a larger number of developed models, can be explained by the dominance of state-driven programs that prioritize rapid deployment and experimentation over large-scale commercial monetization. The pace of digital transformation and innovation calls Europe's traditional industrial and manufacturing leadership into question – and places new demands on European companies to level up on speed, flexibility and technological integration. ▶ A

A Notable AI foundation models developed in 2024

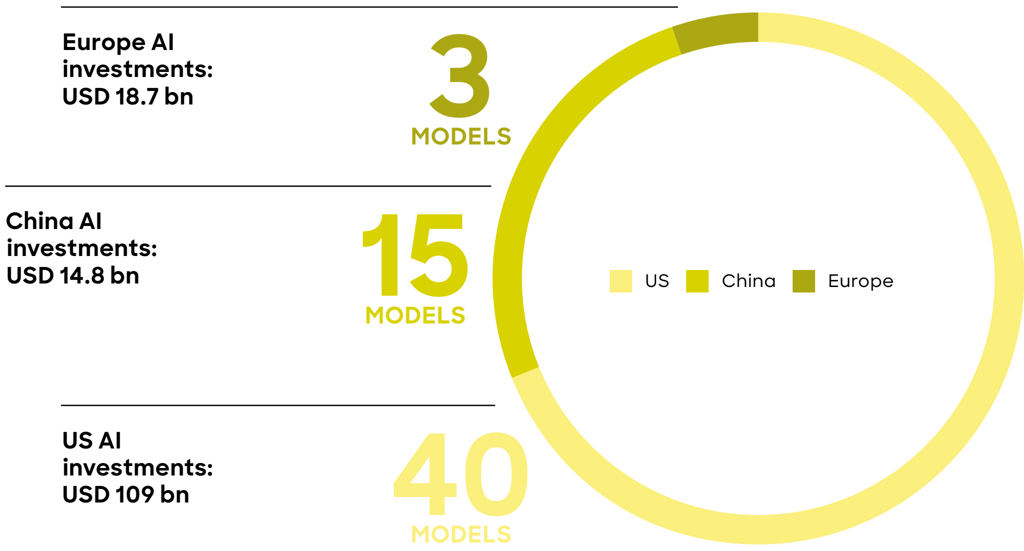

Source: Stanford HAI AI Index 2025, Chapter 1.3

In this evolving landscape, Europe must shift from being a passive technology adopter to an active shaper of its industrial future. European industrial sovereignty no longer depends solely on physical supply chains, but increasingly on who controls the digital systems that steer industrial operations. As processes become more data-driven, real-time and adaptive across industries, AI is emerging as the decisive technology layer that governs these value-adding processes. The strategic implication is as simple as it is far-reaching: **sovereignty over AI systems is not optional – it is the foundation of industrial competitiveness.**

Independence is the key concept here. Sovereignty over AI enables companies to retain ownership and control over their proprietary industrial data. It also allows them to adapt AI solutions to domain-specific requirements across safety-critical and regulated processes. It maintains flexibility in system evolution and supplier choice. And it means they can avoid technological, data, operational and – perhaps most importantly – value lock-in. AI sovereignty isn't about isolation – it's about maintaining full control over the value creation engines that will define the next era of industrial leadership.

B EU AI Act timeline

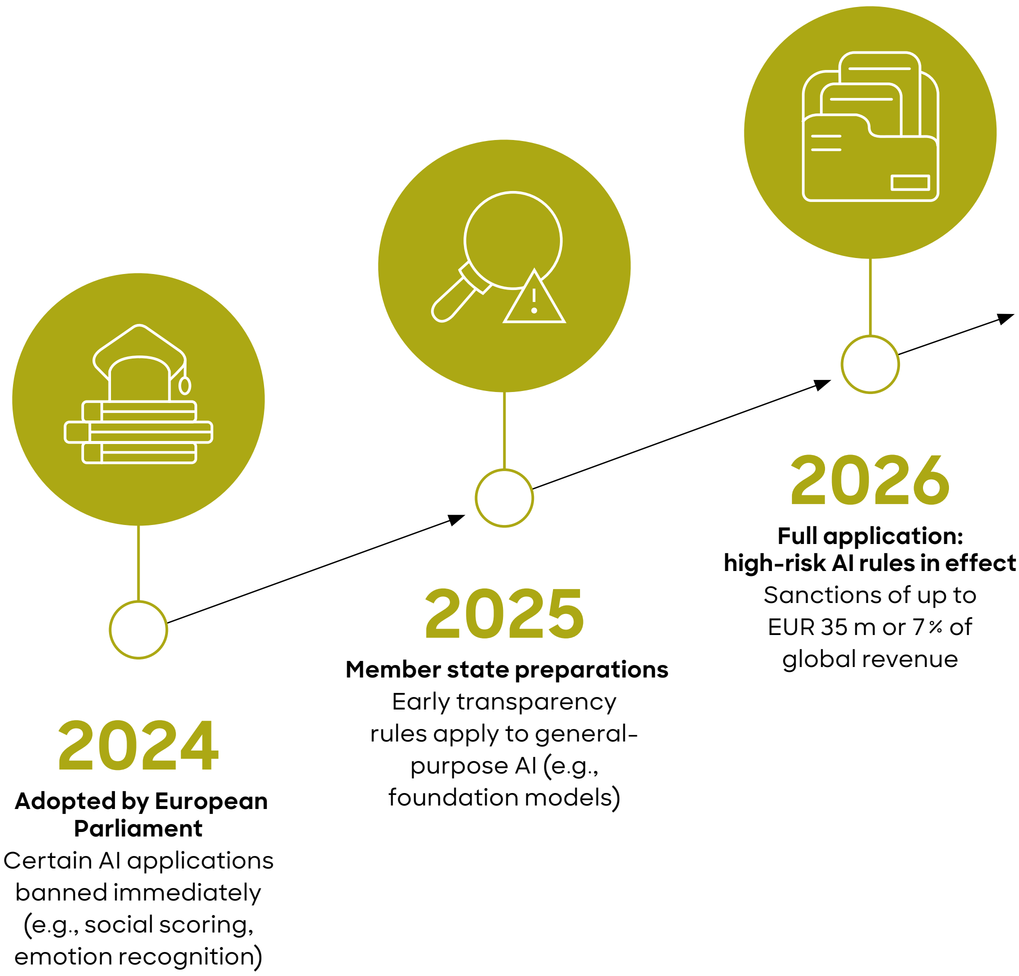

Source: EU Regulation 2024/1689 – EU AI Act

Of course, this is not news for policymakers in the EU. The EU AI Act – the first major legal framework in the world specifically designed to regulate AI systems – was adopted by the European Parliament in 2024 and is expected to be fully implemented by 2026, with some rules applying earlier. The Act defines a risk-based regime with strict requirements for "high-risk" AI systems – including any industrial AI system, especially those highly regulated sectors such as manufacturing, energy, transportation or life sciences with sensitive trade secrets that have automation, predictive maintenance, quality control or safety-critical processes. In the automotive sector, for instance, original equipment manufacturers (OEMs) and Tier 1 manufacturers must demonstrate audited transparency, robust risk management, technical documentation and post-deployment monitoring starting in 2026. Failure to comply may result in fines of up to EUR 35 million or 7% of global revenue. ▶ B

As part of InvestAI, a complementary initiative, the EU has pledged EUR 200 billion to build a domestic AI ecosystem – with local compute, data access and regulatory infrastructure. Work is underway on 13 "AI factories" based on high-performance computing, with plans for three to five large-scale "AI gigafactories" by 2027, each expected to house around 100,000 graphics processing units. The goal is to triple Europe's compute capacity for training and deploying advanced AI models.

Initiatives such as these are attempts to close the EU's massive investment gap compared with the US – estimated at USD 1.36 trillion just in cloud companies and information and communications technology companies. Whether its latest initiatives will erase this structural disadvantage and restore European industry to world-class competitiveness remains uncertain. One thing remains clear: in today's global economy, AI sovereignty isn't a luxury – it's a strategic necessity for European industry.

2

Europe's AI playbook

What sovereignty requires

The EU's strategic response to external dependencies and internal innovation pressure lies not in imitation, but in differentiation. The goal should be to build AI systems that reflect European industrial values, regulatory frameworks and sector complexity.

The good news is that Europe has a number of unique advantages here. The EU is home to a globally leading industrial base, accounting for roughly one-quarter of global manufacturing output. For example, it is the largest exporter of machinery and transportation equipment globally, with business valued at over EUR 1.2 trillion annually. Its industry has rich domain expertise, high-quality engineering data and complex manufacturing ecosystems. What is more, Europe also has long-standing regulatory capabilities and standardization cultures that enable AI to be developed in a trusted, auditable and accountable manner.

Sovereign AI must deliver on four core requirements. First, it must embed trust by design – not by promise – through auditable models, explainability and full compliance with the EU AI Act. Second, control over core assets is essential: data must remain with the company;

> "When you rely on foreign platforms for core AI, you're not just outsourcing tech – you're giving up your data, your value and your edge."
>
> **Daniel Rohrhirsch, Senior Partner, Roland Berger**

deployment must support on-premise, private cloud or secure shared cloud environments; model governance must be transparent; and integration of AI must be modular. Third, domain-specific needs take precedence; customized solutions that embed deep domain knowledge consistently outperform generic models. Lastly, sovereign AI must scale compatibly with Europe's existing infrastructure, integrating with legacy systems and supporting edge or cloud hybrid architectures. ▶ C

While the US and China race ahead on innovation speed, Europe's AI strategy plays a different game. The EU prioritizes trust, transparency and industrial customization, all built on a foundation of clear regulation and standardization – embedding sovereignty not just as a political principle but as an operational advantage. In the US, AI models are owned by hyperscalers and data is controlled by platforms; in China, the state coordinates

C Four requirements of sovereign AI

Trust by design Auditability, explainability, EU AI Act compliance	**Modular control** Data sovereignty, on-premise deployment, modular integration
Domain fit Domain-specific AI, tailored to regulated industries	**Infrastructure fit** Legacy system integration, hybrid edge/cloud scaling

Source: Aleph Alpha

D **Reference framework: Global AI models**

Dimensions	US (hyperscaler-led)	China (state-led)	EU (sovereignty-led)
Innovation speed	Very high (scale-driven)	High (state-coordinated)	Medium (fragmented)
Data control	Platform-owned	State-access	User-or company-owned
Model governance	IP-driven	State-prioritized	Transparent/ auditable
Trust level	Low	Low	Very high
Industry customization	Low	Medium	Very high (domain-focused)

Source: Roland Berger

development and maintains broad access to training inputs. In contrast, Europe's emerging AI stack emphasizes user-owned data, auditability, explainability and domain-specific tuning. The EU AI Act codifies this approach, mandating transparent and accountable models across industrial use cases. Rather than chase hyperscaler speed, Europe's AI sovereignty strategy carves out a unique competitive position: enabling compliant, controllable and customizable AI tailored to its advanced industrial base. ▶ D

3

From design to factory floor

Where AI delivers an industrial edge

European industrial companies are under growing pressure to accelerate time-to-market. Shorter development cycles are being driven by disruptive technologies, rapidly shifting customer expectations, global competition and ever-tighter regulatory demands. At the core of this transformation lie two interconnected challenges: the need for greater speed and the rising complexity across the entire product development and manufacturing lifecycle.

Three distinct structural challenges are reshaping industrial product development. First is the explosion of requirements. Products must now satisfy an ever-growing array of demands – functional, regulatory, environmental and safety-related. Documentation on recyclability and carbon footprint is no longer optional but integral to the design process.

This creates intricate interdependencies between mechanical, electrical, electronic and software components, including embedded systems and Internet of Things (IoT) capabilities.

The second major challenge is the surge in data volumes and the complexity of data management. Complex assemblies for computer-aided design, multiphysics simulations, verification and validation reports, and test and measurement data are driving exponential growth. Managing this dynamic information – especially as designs evolve and regulatory expectations shift – requires responsive data processes and real-time synchronization. Ensuring consistency and traceability of heterogeneous data types across product lifecycle management (PLM) systems and model-based systems engineering (MBSE) approaches is increasingly difficult.

Third, interdisciplinary collaboration is becoming essential, requiring role-specific tools and seamless workflows. Engineers, developers and domain experts can no longer operate in isolated silos. Their methods and toolchains must be integrated into coherent development processes. To support this, AI-powered assistants and digital tools must deliver user-specific interfaces and dashboards that provide relevant, actionable information.

These three forces can be found along the different phases of development processes. As a result, achieving efficiency, quality assurance and timely market delivery has become significantly more challenging. Using the conventional V-model offers a good opportunity to visualize these different sub-processes.

AI proves especially valuable in product development by breaking the process into modular, well-defined sub-processes that can be individually enhanced through specialized AI solutions. This enables five key use cases where AI delivers significant impact: requirements engineering, architectural design, product validation, release management and systems traceability. These five key use cases are discussed in turn below. ▶E

" AI for industry has to work with what's already in place: legacy systems, hybrid infrastructure and shifting regulations."

Patrick Ruthardt, Principal, Roland Berger

E The value added by Pharia Engineering

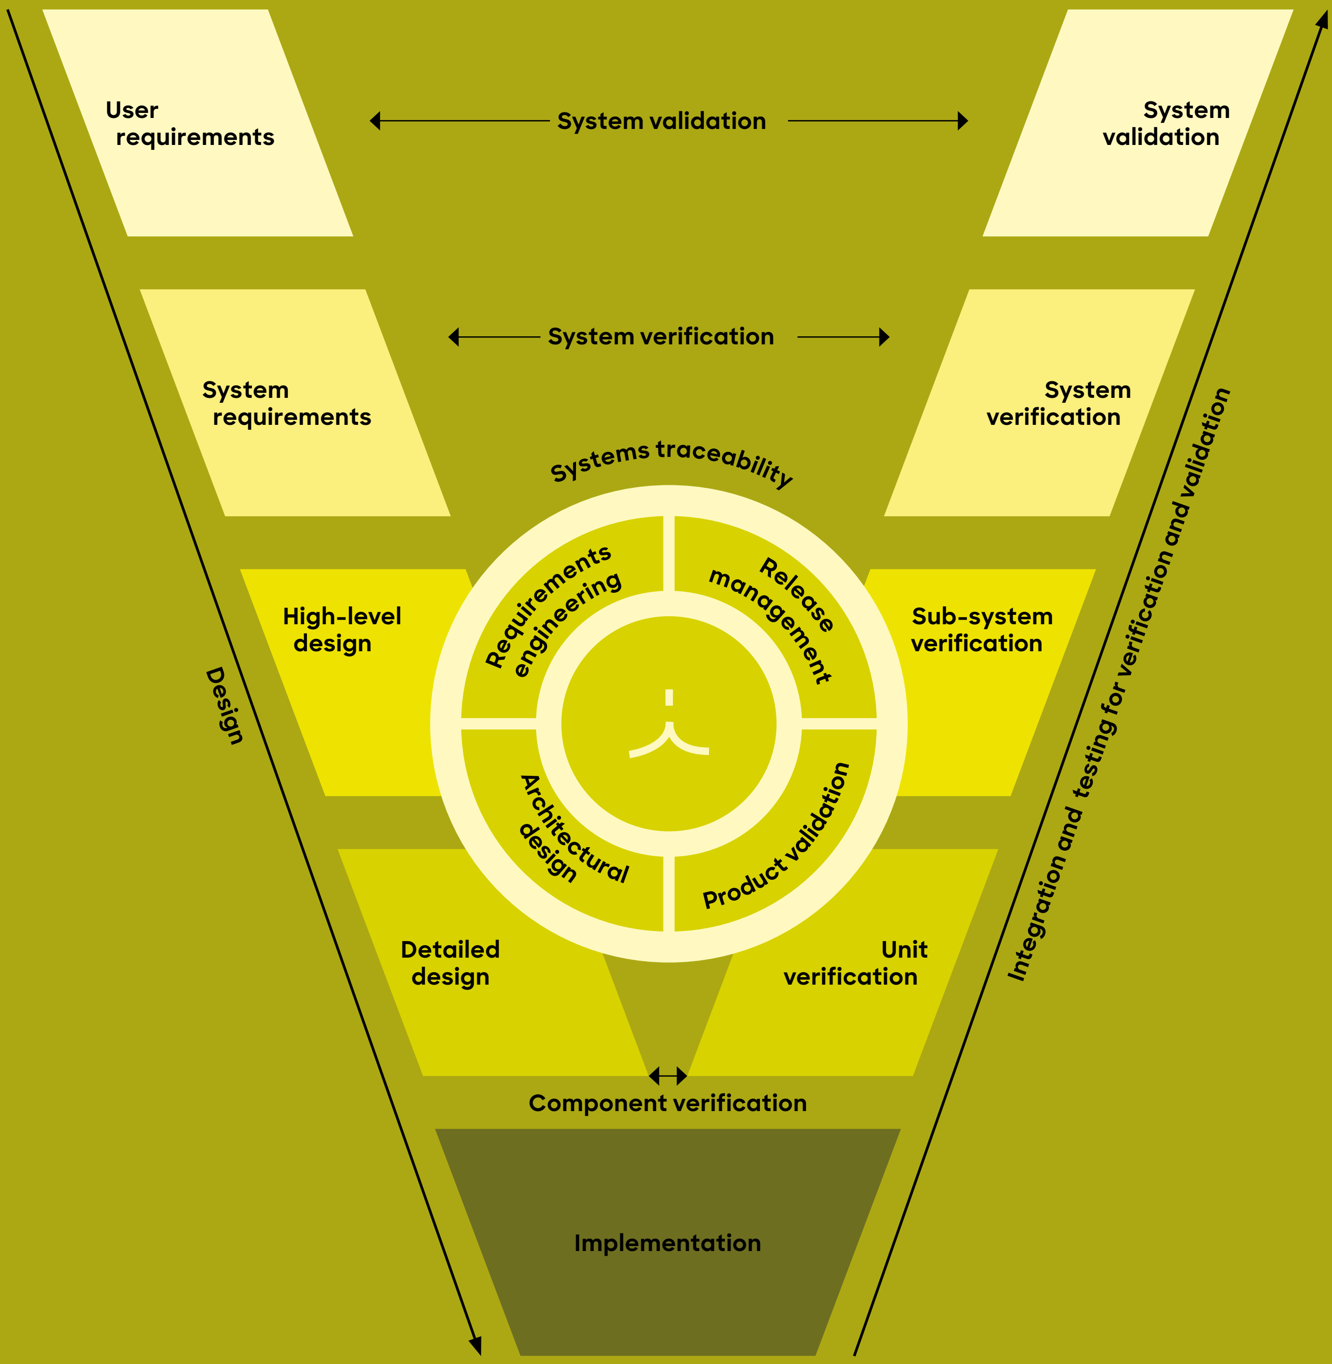

Source: Roland Berger

#1 REQUIREMENTS ENGINEERING

The requirements engineering process begins with an analysis of the request for quotation (RFQ) documents, extracting and structuring customer requirements. These redundancies are then evaluated, categorized and assigned to the appropriate engineering teams. Gaps and inconsistencies are identified and addressed through iterative discussions with the client. A compliance matrix is created to ensure full coverage of all requirements. As product development progresses, detailed requirements are mapped to the system architecture and undergo ongoing review – though, in practice, this ideal is often hindered by several persistent challenges.

Large, unstructured RFQ packages often arrive in diverse formats such as PDFs, Excel sheets and images, demanding high manual effort to extract, translate and document requirements. Additionally, past project knowledge is often underutilized, limiting its potential to close gaps or assess risks. These requirement reviews are time-consuming, with repeated iterations needed to ensure completeness, consistency and risk mitigation.

To address these inefficiencies, an AI-powered requirements engineering assistant offers targeted support across all stages of the requirements engineering process. During RFQ processing, AI applies semantic extraction to automatically structure data from incoming documents – accelerating initial data ingestion and significantly reducing manual workload. AI helps categorize, translate and map requirements using databases and ontologies, and highlights differences from prior projects. It also supports the generation of compliance matrices by automating completeness checks and helping produce regulatory coverage reports.

#2 ARCHITECTURAL DESIGN

Architecture activities transform customer requirements into functional and interface definitions at both the system and sub-system levels. The process starts with high-level architecture derivation, followed by sub-system detailing. Architecture refinement combines existing building blocks, explores multiple architectural variants and involves consistency checks and compliance validation. At the sub-system level, functions are mapped to components, technical feasibility is assessed and interface consistency is verified.

This process faces challenges due to the complexity of multi-layer architectures, where interdependencies across systems, sub-systems and interfaces can be overwhelming. Teams often lack a comprehensive understanding of system constraints, including physical, technical and regulatory limitations.

To help manage this complexity, AI provides support across key architectural activities. It assists in automatically matching functions to requirements by leveraging historical project data and domain expertise, enabling the rapid generation of multiple architecture variants for early-stage exploration. AI helps identify interface inconsistencies, gaps or conflicts early on and supports validation against various constraints, such as regulatory standards or physical limitations.

#3 PRODUCT VALIDATION

The process begins by decomposing the system into components and functions, identifying potential failure modes, root causes and consequences. Interdependencies are mapped to understand failure propagation. Risks are evaluated using severity, occurrence and detection metrics, forming a risk priority number. This enables targeted countermeasures. Design validation (DV) confirms that the design meets requirements; product validation (PV) ultimately ensures that it meets customer needs under real conditions.

DV/PV insights feed back into the design-failure mode and effects analysis (or D-FMEA) to refine and strengthen the design. Nevertheless, several challenges continue to undermine the robustness of this process. Critical failure modes may go undetected due to incomplete system decomposition, overlooked interdependencies or a lack of cross-functional knowledge. In many cases, test plans become oversized, a result of conservative strategies that assign equal weight to all scenarios rather than prioritizing based on actual risk levels. The manual workload involved in analyzing failures, assessing risks, documenting findings and aligning test plans across multiple stakeholders remains high.

To alleviate these burdens, AI provides targeted support throughout key process stages. It identifies potential failure modes by analyzing historical data and FMEA reports, supports system decomposition and maps functional dependencies. For risk evaluation, it leverages past project data to score severity, occurrence and detection more accurately. It accelerates the development of countermeasures by suggesting proven design solutions and optimizes test plans by eliminating redundant tests. It also enhances testing by focusing resources on high-risk areas and minimizing unnecessary validation efforts. Finally, it ensures continuous FMEA updates by integrating test results and field data.

#4 RELEASE MANAGEMENT

The release management process verifies system and software requirements across model, software, hardware and vehicle-in-the-loop environments. It begins by defining a test profile based on architecture and requirements, followed by detailed planning of test types, test cases, expected outcomes and configurations. Execution involves running tests, collecting logs and evaluating results against criteria. A refinement phase updates test cases based on defect analysis and regression feedback, ensuring earlier defect detection and more reliable system performance.

Still, this process can be slowed by inconsistent, manual documentation and insufficient test coverage, particularly in software-in-the-loop environments. Large data volumes make defect analysis and coverage tuning increasingly complex and resource-intensive. Test case reuse remains limited, while fragmented collaboration between designers, engineers and developers complicates alignment. Integration across toolchains is often weak, further reducing process efficiency and making it difficult to maintain momentum in iterative development cycles.

AI addresses these issues by generating consistent test cases from requirements and historical data, optimizing plans through defect pattern analysis and converting documentation into machine-readable formats. It automates test report and log analysis, enabling faster defect detection and correlation with code changes. AI also recommends validated test cases and configurations from earlier cycles, improving reuse, consistency and integration with continuous development and CI workflows.

#5 SYSTEMS TRACEABILITY

A centralized systems traceability capability is at the core of engineering efficiency, enabling fragmented experience to be leveraged as accessible knowledge. It starts with semantic search across internal and external sources, retrieving and prioritizing relevant insights. These are organized into actionable assets that engineers can apply directly. AI then generates draft documents, reports or templates, integrating standards or expert input as needed. Users iteratively refine the content, enhancing quality and relevance. Continuous updates from projects and structured sources keep the hub current and evolving.

However, in many organizations, traditional knowledge-sharing methods are slow and inefficient. Engineering know-how is scattered across tools, teams and documents, making retrieval difficult. Content creation remains largely manual, and duplication of effort is common due to poor reuse. Organizational silos disconnect related information such as requirements, risks and tests, while isolated data limits visibility into project history, reducing the quality of decision-making and slowing design maturity.

AI addresses these issues by creating a unified platform that integrates knowledge across engineering domains. The key differentiator from hyperscaler enterprise generative pre-trained transformers is that, with flexible hosting options, even highly sensitive engineering or domain-specific data can be securely connected. Semantic search enables users to retrieve meaningful, cross-functional insights rather than isolated files. AI links related data such as requirements, failure modes and test cases to reveal hidden connections. It delivers context-aware answers based on complete project history and helps teams reapply lessons learned. This accelerates content creation, breaks down silos and strengthens decision-making across engineering teams.

CASE STUDY

Case study 1
Boosting efficiency on the shopfloor

Industrial production is under increasing pressure: rising costs, especially in high-wage regions, combined with more-complex machinery and growing automation, are straining factory operations. Meanwhile, skilled labor shortages further compound the challenge. AI offers new solutions to improve overall equipment effectiveness (OEE) by delivering tailored insights and recommendations to both frontline workers (operators, technicians) and managerial staff.AI applications in factories align with different layers of the automation pyramid and serve various roles – maintenance technicians, machine operators and production managers alike. While traditional Industry 4.0 tools (such as predictive maintenance) focus on sensor data and the IoT, generative AI (GenAI) now also unlocks value from unstructured data such as service manuals, logs and maintenance tickets.

Use case 1:
AI-customized maintenance instructions

Maintenance tasks are often hampered by generic instructions and limited troubleshooting support, leading to significant downtime costs. Technicians are left to interpret documentation while unstructured data from service tickets remains underutilized. GenAI changes this by generating interactive, step-by-step maintenance guides tailored to the issue at hand, drawing on manuals, historical service cases and contextual insights. The result is faster diagnosis, quicker resolution and reduced equipment downtime.

Use case 2:
AI-powered key performance indicator (KPI) assistant

Production managers frequently struggle to access performance data due to siloed and complex information technology/operational technology (IT/OT) systems. Transparency is limited and navigating these platforms requires time and technical know-how. A GenAI assistant allows users to ask questions in natural language, pulling data from across systems (such as supervisory control and data acquisition, manufacturing execution systems, PLM and enterprise resource planning) to build personalized, visual reports and provide instant answers on performance metrics. This improves visibility and enables smarter, faster decision-making, ultimately supporting long-term OEE improvements.

Use case 3:
Prescriptive maintenance (machine learning and GenAI)

Standard predictive maintenance systems rely mainly on time-series data and often lack insight into why failures occur. Instructions are not customized, and there is only limited automation of maintenance tasks. GenAI can solve this challenge by combining sensor data with unstructured text analysis to identify root causes and explain them. It then generates targeted repair instructions and can automate subsequent maintenance tasks. This reduces downtime, boosts OEE and lays the groundwork for semi-autonomous or self-maintaining production systems – building directly on the capabilities developed in the previous two use cases.

> **"Predictive maintenance helps identify problems early. But, it often falls short of explaining their root causes or suggesting how to fix them."**
>
> **Jan Stratmann, Solutions Principal, Aleph Alpha**

Why do European AI models fit so well with European industry? Because they are purpose-built to meet the region's unique demands for compliance, security and deep customization across highly regulated and technically complex sectors. These models ensure full data sovereignty, safeguarding industrial IP and sensitive customer data – an essential requirement for manufacturers, automotive companies, pharmaceutical firms and their suppliers. This enables them to deploy AI without compromising proprietary know-how. Their transparency and auditability support strict regulatory compliance, a necessity in domains like healthcare, automotive safety and financial services. Crucially, European AI models are designed for deep industrial customization, aligning closely with the nuances of manufacturing processes, engineering workflows, supply chains and industry-specific standards. Transparent model governance facilitates certification and traceability, reinforcing quality control and legal accountability. Finally, their alignment with the EU innovation model supports decentralized supplier ecosystems, allowing both large OEMs and small and medium-sized enterprises (SMEs) to adopt AI collaboratively – without the risk of platform lock-in or dependence on non-European cloud providers.

Building an AI-driven engineering hub requires a strategic, phased approach rooted in both technical depth and domain relevance. It begins with structuring engineering knowledge from sources such as requirements, architectures and validation reports. This enables domain-specific AI to ensure accuracy. The initial focus should be on automating requirement documentation, with later expansion into architecture design and consistency checks. Modular design is key, allowing models to scale across testing, FMEA and validation planning. The goal is not just automation, but intelligent orchestration of engineering knowledge.

A final – but essential – point: while AI serves as a powerful accelerator and assistant, it does not replace expert judgment. Engineering teams retain full responsibility for validation, critical decision-making and the creative problem-solving that drives innovation. AI enhances the process; it does not substitute the expertise. ▶F

F Step-by-step expansion of the Pharia Engineering suite toward domain intelligence

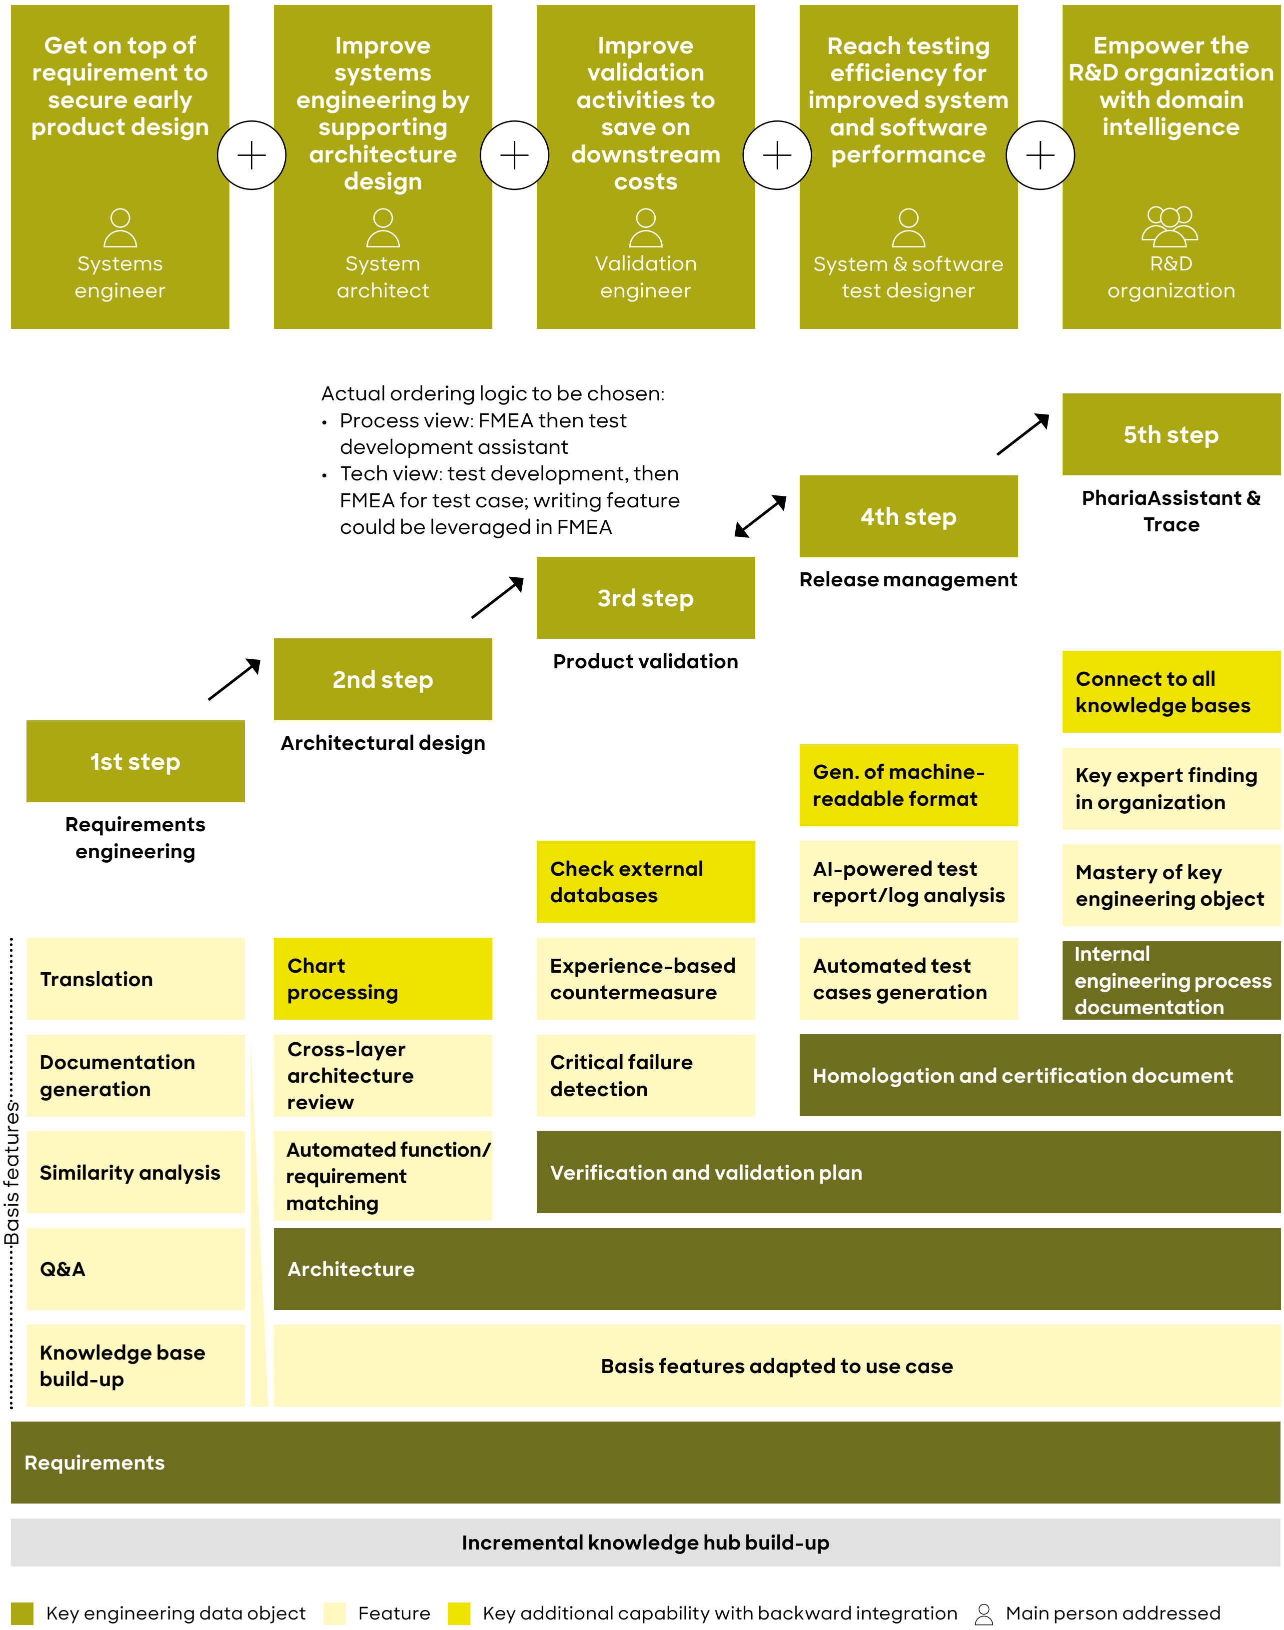

Source: Roland Berger, Aleph Alpha

4

From principles to practice

What it takes to build sovereign AI

Turning the vision of sovereign AI into reality calls for more than abstract principles – it requires concrete capabilities purpose-built for industrial needs. The sovereignty question is not limited to geopolitics or whether the EU or the US will lead as the origin and home of AI technology; it also extends to the degree to which individual companies can maintain control over their own sovereign AI systems and retain ownership of their data.

G Making sovereign AI a reality

Source: Roland Berger

"In regulated sectors, trust starts with transparency, version control and keeping data on-site from the start."

Andreas Geiss, VP Industrial & Retail, Aleph Alpha

At the core of this challenge are four key enablers: transparent data and models, hybrid edge/cloud infrastructure, customization and modularity, and regulatory compliance by design, particularly with regard to the EU AI Act. These elements ensure that sovereign AI is not only effective but also auditable, secure and aligned with the complex demands of regulated industries. ▶ G

To meet these needs, purpose-built AI suites are emerging – such as PhariaAI from Aleph Alpha – that offer modular, secure and controllable deployment environments. These AI suites support both cloud-based and on-premise or air-gapped operations, making them adaptable to a broad range of operational and compliance contexts, including sensitive sectors such as defense, healthcare and critical infrastructure.

A defining feature of these systems is a unified AI control layer that acts as a bridge between models, data and applications. This control plane allows organizations to safely manage customization, maintain detailed logs, enforce governance policies and meet audit requirements from a single interface.

Crucially, these architectures aim to be trustworthy by design. They combine explainability, traceability and modular control in a coherent framework, ensuring that industrial AI is not only powerful but also governable – in other words, ready for real-world deployment in environments where trust, control and compliance are non-negotiable.

5

From pilot to impact

How to implement industrial AI

In industrial settings, one of the clearest success factors for effective AI implementation is the ability to combine and contextualize data across the full automation pyramid, from IT systems to OT on the ground. This integration is where AI really comes into its own, bridging silos and enabling more responsive, data-driven operations. That said, industrial environments are still heavily dependent on time-series data and machine vision, meaning traditional machine learning remains essential. GenAI doesn't replace these, it builds on them. A highly effective starting point is already available in the form of GenAI-powered, customized maintenance instructions, which are helping technicians reduce downtime and improve task accuracy today. Looking ahead, the next frontier lies in prescriptive and semi-autonomous maintenance and operational systems, which will rely on tighter integration between GenAI and time-series-based IoT applications. Throughout this evolution, it's essential to understand the business impact of each use case, ensuring that AI efforts are grounded in real operational value and prioritized accordingly.

In R&D and product development, the benefits of AI are equally compelling. Organizations are seeing efficiency gains of 20% to 40%, driven by automation in analysis, review and documentation workflows. But the value extends beyond speed: AI also enhances design quality by supporting consistent, experience-based decisions. By supporting requirement formulation, architectural consistency and validation planning, AI enables early risk mitigation across the product development cycle. This contributes to higher product maturity and robustness, as critical engineering steps are reinforced with AI-generated insights. The result is a faster, more reliable path from concept to market.

The following two case studies illustrate how organizations are putting these principles into action – moving beyond pilots to scaled, value-generating AI systems that align with their engineering workflows and business goals.

"AI sovereignty isn't about putting up barriers. It's about giving European industry the tools to lead with confidence."

Philipp Frey, Project Manager, Roland Berger

CASE STUDY

Case study 2
A scalable GenAI assistant for industrial maintenance

A global industrial technology leader faced growing complexity in maintenance operations. Engineers relied on multiple tools – IoT dashboards, ticketing systems, documentation portals and handwritten logs – to resolve issues. These fragmented workflows led to delays, knowledge silos and reliance on experienced staff. Tools lacked personalization and intuitive access, increasing cognitive load and reducing task efficiency. Persistent challenges included downtime, inefficient shift handovers and knowledge loss during turnover.

To address this, Aleph Alpha deployed PhariaAssistant and Trace, a GenAI-powered assistant, across global production sites. Technicians interact via text or voice to receive precise, context-aware answers based on real-time sensor data, service histories and documentation. The system ensures transparency, traceability and secure, on-premise deployment aligned with enterprise compliance standards.

Key capabilities include:

Smart maintenance assistant

- Generates repair guides tailored to machine type, failure history and technician language
- Fuses historical tickets, sensor data and documentation into situational answers
- Includes source links for every step, reducing resolution time by up to 25%

Audit and KPI assistant

- Allows natural language queries of uptime, energy use and quality metric
- Auto-generates charts and reports for shift logs and audits
- Supports hands-free speech-to-text input, improving OEE by up to 8%

Prescriptive maintenance and root cause analysis

- Combines time-series and unstructured data to explain failure causes
- Outputs interpretable, citation-based diagnostics
- Suggests spare parts, follow-ups and draft service tickets

Following a successful pilot, a multi-site rollout was initiated. Early results showed faster troubleshooting, improved knowledge reuse, and smoother shift transitions. The initiative is now expanding into additional use cases – including energy, safety and training – with full IT/OT integration planned.

CASE STUDY

Case study 3 Agentic AI for high-fidelity industrial research

A leading industrial enterprise sought to accelerate complex R&D and market research workflows without compromising quality or compliance. Teams operated across domains – engineering, strategy, compliance – relying on scattered sources such as internal reports, patents, regulations and market data. Manual research cycles spanned weeks, constrained by siloed expertise and inconsistent tools. Speed, traceability and multimodal input handling became critical, especially under strict EU data governance rules that ruled out opaque, cloud-only AI solutions.

Aleph Alpha deployed an agent-based Deep Research System, designed to augment, not replace, domain experts. Built on sovereign open-source large language models, it operates in secure hybrid or on-premise environments. The system orchestrates over 30 specialized AI agents that work in parallel, mimicking expert workflows: extracting requirements, benchmarking designs, scanning IP and validating compliance. Every output includes confidence scores, source links and relevance indicators to ensure trust and transparency.

Key capabilities include:

Long-form research acceleration

- Generates 20- to 50-page reports from diverse inputs in hours, not weeks
- Includes citations, summaries and SME-level analysis across technical, market and legal dimensions

Patent and risk landscape analysis

- Identifies IP gaps, innovation clusters and key players using live patent data
- Provides traceable risk profiles and scenario analyses for strategy and R&D teams

Integration with the V-model lifecycle

- Supports requirement extraction, system design references, implementation roadmaps, QA checks and trend monitoring
- Enables continuous insight generation from pre-design to post-market phases

The deployment delivered time savings of more than 80 % in research cycles and significantly improved cross-department collaboration. Analysts now rely on traceable, repeatable insights rather than institutional memory. The system laid the foundation for hybrid research teams combining human oversight with autonomous AI agents – offering a repeatable model for research excellence in regulated industries.

Recommendations

We recommend five clear steps for industrial leaders to implement AI effectively and securely. First, define your AI governance and control requirements – determine the degree of oversight, compliance alignment and data autonomy your organization needs to operate safely and confidently. Then, prioritize high-impact AI use cases that solve real problems and deliver measurable value, such as predictive maintenance, automated documentation or AI-driven quality control.

Next, build safeguards for auditability and governance from the start. Ensure transparency, traceability and compliance readiness, especially under evolving regulations such as the EU AI Act. At the same time, integrate AI into your OT or IT systems without lock-in. Choose modular, interoperable solutions that avoid vendor dependency and work across existing infrastructure.

Finally, partner strategically – combine domain-specific AI with proven implementation expertise. The right collaboration ensures that solutions are not only technically sound but also deeply aligned with real-world operational needs.

This white paper was co-authored by experts from Roland Berger and Aleph Alpha, combining strategic consulting insights with cutting-edge AI expertise.

Daniel Rohrhirsch
Senior Partner
daniel.rohrhirsch@rolandberger.com

Patrick Ruthardt
Principal
patrick.ruthardt@rolandberger.com

Philipp Frey
Project Manager
philipp.frey@rolandberger.com

Andreas Geiss
VP Industrial & Retail
andreas.geiss@aleph-alpha.com

Jan Stratmann
Solutions Principal
jan.stratmann@aleph-alpha.com

Suraj Shetty
Solutions Engineer
suraj.shetty@aleph-alpha.com

CONTRIBUTORS

Bernhard Langefeld
bernhard.langefeld@rolandberger.com

Denis Depoux
denis.depoux@rolandberger.com

Felix Mogge
felix.mogge@rolandberger.com

Florian Deichmueller
florian.deichmueller@rolandberger.com

Hugo Carreira
hugo.carreira@rolandberger.com

Manuel Schieler
manuel.schieler@rolandberger.com

Marcus Hartmann
marcus.hartmann@rolandberger.com

Maria Mikhaylenko
maria.mikhaylenko@rolandberger.com

Michael Schich
michael.schich@rolandberger.com

Nizar Hneini
nizar.hneini@rolandberger.com

Sebastian Feldmann
sebastian.feldmann@rolandberger.com

Further reading

- *BEYOND AUTOMATION: WHY AI AGENTS ARE YOUR NEXT STRATEGIC IMPERATIVE*
- *GLOBAL AUTOMOTIVE SUPPLIER STUDY 2025*
- *ARTIFICIAL INTELLIGENCE*

09.2025

ROLANDBERGER.COM

ROLAND BERGER is one of the world's leading strategy consultancies with a wide-ranging service portfolio for all relevant industries and business functions. Founded in 1967, Roland Berger is headquartered in Munich. Renowned for its expertise in transformation, innovation across all industries and performance improvement, the consultancy has set itself the goal of embedding sustainability in all its projects. Roland Berger generated revenues of around 1 billion euros in 2024.

ALEPH ALPHA was founded in 2019 with the mission to research and build sovereign, human-centric AI for a better world. With a team of international scientists, engineers, and innovators, the company researches and builds transformative AI solutions. Its full stack offering, PhariaAI, is the prime choice for enterprises and government institutions seeking to retain sovereignty, secure their data, and build trustworthy applications.

Publisher

Roland Berger GmbH
Sederanger 1
80538 Munich
Germany
+49 89 9230-0

25_2210_REP