rescale

AI-Powered Engineering and the Future of R&D

How domain-specific advances in agentic digital engineering and AI physics are reigniting R&D productivity and accelerating new product development

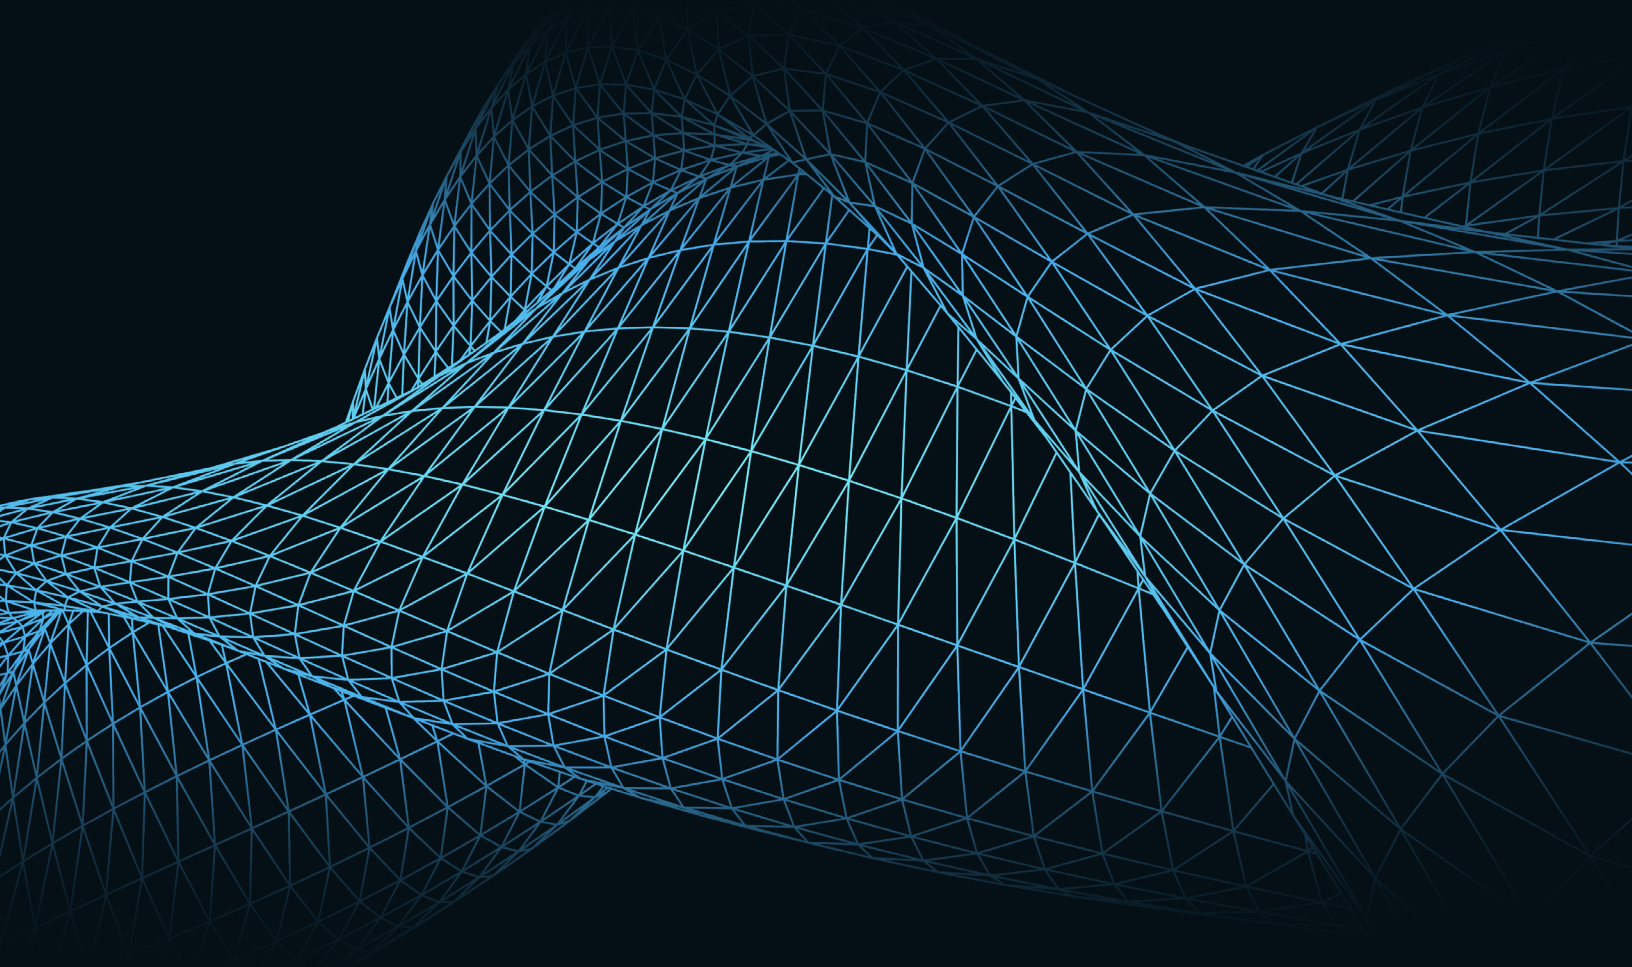

Table of Contents

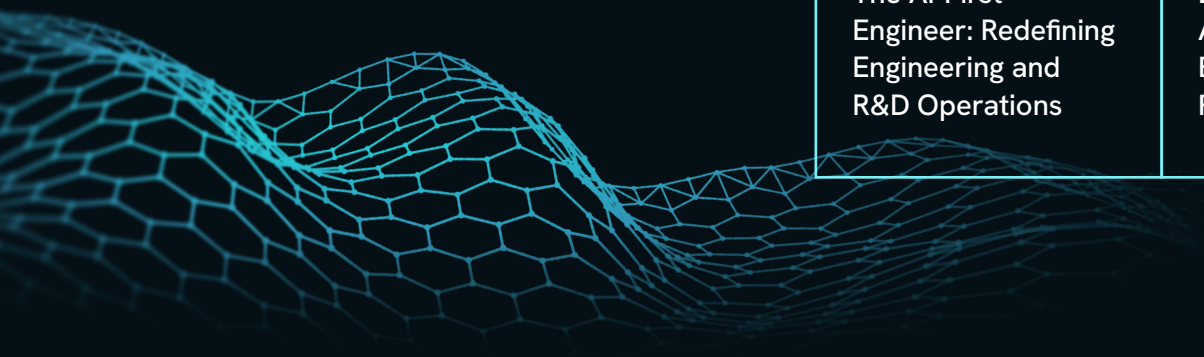

Reigniting R&D Productivity With AI

Innovation has been the engine of human progress for two centuries. AI-powered R&D is poised to unlock a new era of breakthroughs.

For most of the modern era, the relationship between R&D investment and innovation output was broadly reliable: more investment yielded faster breakthroughs. Today, that paradigm doesn't always hold true as many sectors have faced decreasing research efficiencies.

Simply put, most industries are spending more for less innovation. McKinsey research on R&D innovation across sectors documents a pervasive and persistent challenge: R&D productivity is declining across industries.

- In semiconductors, maintaining Moore's Law required 18X more inflation-adjusted R&D spending now versus the 1970s.[1]
- In pharmaceuticals, the number of new drugs approved per billion dollars of R&D spent reduced by 50% roughly every nine years between the 1950s and today — a phenomenon researchers dubbed 'Eroom's Law,' the reverse of Moore's.[1]
- Across materials science and other engineering disciplines, the same pattern emerges: each dollar spent on R&D buys less innovation than it once did.

> "Engineering teams today need a platform that integrates computational engineering, data intelligence, and AI into a single environment where every workflow builds on the last, continuously turning R&D expertise into organizational intelligence."
>
> **Joris Poort,**
> Founder and CEO, Rescale

18X More R&D spend required to sustain Moore's Law[1]

50% Decline in drug discovery output per R&D dollar every 9 years[1]

$500B+ Annual value AI could unlock through accelerated R&D[1]

The root cause is not a lack of talent or ambition. It is the relentless growth in complexity — of product requirements, of data volume, of regulatory requirements — that has been outpacing the toolsets available to engineering teams. Scientists and engineers are simply spending too much of their time managing processes and legacy technologies rather than creating.

Yet scientific and engineering breakthroughs remain the greatest catalyst for economic growth and social wellbeing. And there is a clear path emerging to reverse R&D's current trajectory — rethinking and reengineering R&D around an AI-first operating model.

AI applications are ideally suited for managing the increases in data and information volume that have inhibited innovation. As BCG research projects, "AI in R&D will deliver 10-20% reductions in time to market, up to 20% lower R&D costs"[2], and — critically for scientific and engineering centric industries like aerospace, automotive, life sciences and material sciences a path to market leadership for companies that re-architect their operating models around AI.

The Three Pillars of AI-Powered Engineering

AI Powered Engineering represents a mindset, operational and technology shift. You don't just switch it on. But you can start building the right foundation today.

Every major R&D and product engineering organization faces fundamental constraints: Limited compute resources; fragmented, unstructured data; and legacy technology that cannot scale linearly with the innovation. As a result, teams struggle with top-down efficiency mandates and the ability to meet the productivity and innovation challenges facing their industries.

The shift to an AI-driven R&D operating model will not happen overnight, but leading companies are moving fast and building the future of R&D and product engineering around three rapidly advancing technology pillars.

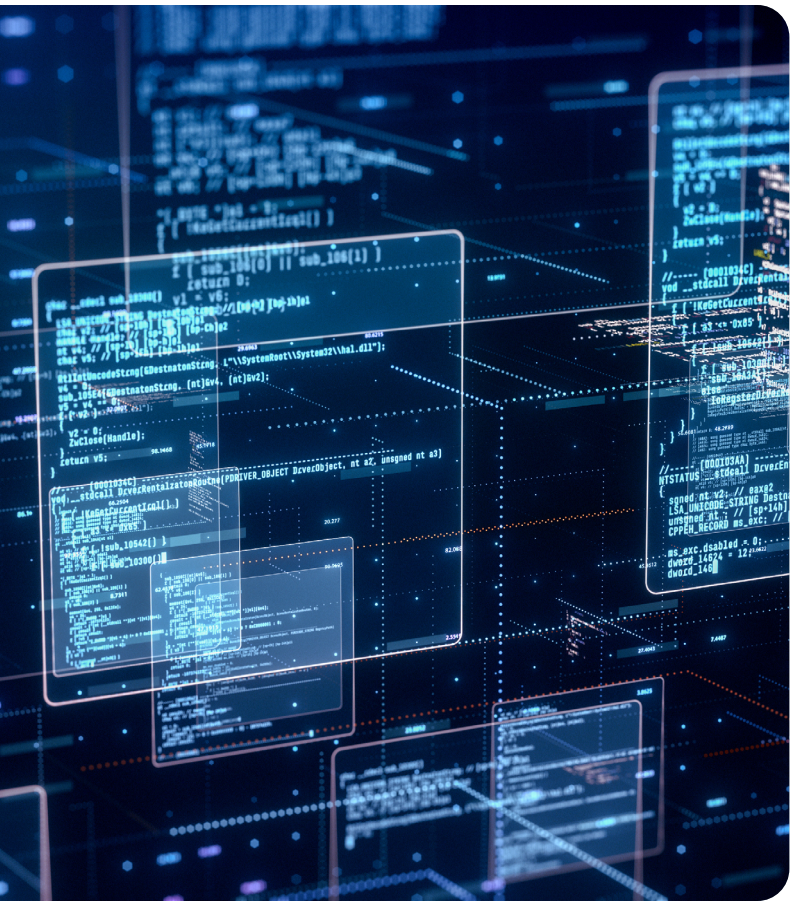

Advanced Simulation: Integrated, AI Assisted and HPC Powered

To accelerate innovation, engineering teams need the ability to run advanced modeling and simulations at the speed of their ideas rather than queue wait times. They need their full simulation stack connected in a single environment with embedded AI applications and on-demand access to high-performance computing.

Unified Product Data: Connected, Automated and AI Ready

A unified data fabric that connects PLMs, CAE tools, simulation archives, and operational systems provides a digital thread that eliminates bottlenecks. Automated tagging, context enrichment, and structuring transforms data into intelligence that engineers can act on faster and serves as a foundation for AI applications.

Engineering AI: Generative, Agentic and Physics AI

R&D and engineering teams need domain-specific AI applications — from generative to agentic to AI physics surrogate models - embedded directly into their engineering workflows. Together this holistic approach to engineering-specific AI, seamlessly integrated into day-to-day workflows becomes a true force multiplier for R&D organizations.

Advancing Simulation: Integrated and Intelligent

While a critical resource, modeling and simulation has not reached its full potential. Disconnected tools and data issues are compounded by computing constraints. It's time to change the model.

Disconnected, Manual and Underpowered

For engineering teams the routine is all too familiar: search for data, manually set up jobs, get in the queue, wait on resources and accept that many exploratory ideas will never be tested. Not because they lack potential, but because resources are constrained.

Engineers are often left to manage complex, disconnected modeling and simulation stacks themselves — solver software licenses; custom data aggregation and structuring; workflow automation scripts; and compute architectures configurations. These heavy manual workloads are compounded by lack of compute resources.

The result is lower fidelity simulations and suppression of engineering creativity and productivity. Current modeling and simulation environments as well as the HPC resources that support them were largely designed around human engineers. As engineering teams begin to leverage agentic engineering and AI Physics to run 24/7 R&D organizations, the current resource constraints will be exacerbated.

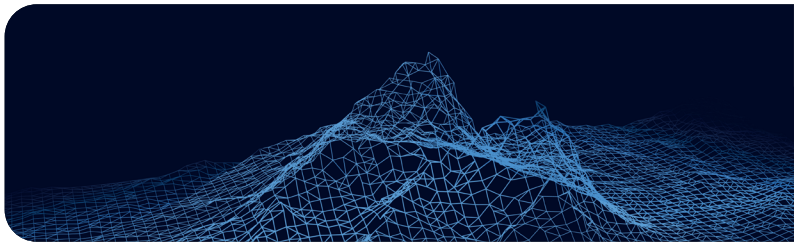

Shifting Forward: Integrated Simulation and Intelligent HPC

Integrated Digital Engineering Platforms have emerged as powerful tools for integrating the full modeling and simulation stack in a single pane experience for engineers — connecting and automating end-to-end modeling and simulation workloads.

- As Digital Engineering Platforms become more powerful through embedded AI assistants, agents and surrogate models they will need more intelligent computing.
- Cloud HPC has been a powerful resource, providing the elastic compute engineers needed on demand, but teams still struggle with performance vs. budget tradeoffs.
- Intelligent HPC driven by invisible orchestration engines will evaluate priority, provider, region, hardware tier, and cost-model for every job and select the optimal configuration.
- And a new breed of optimization tools will help engineering and IT leaders dynamically adjust set ups to balance productivity and speed with hardware and software costs.

This gives organizations the controls they need to adopt cloud-based modeling and simulation environments without runaway cost concerns as engineering and IT leaders can now actively optimize for performance and cost to meet their product development goals.

A Continuous Discovery and Design Engine

The power of integration and intelligent computing is that modeling and simulation transforms from a selective validation tool into an infinite discovery engine. Engineering teams can cost effectively run continuous, automated design space exploration — generating and evaluating millions of variations across multiple physics domains and using their simulation data to train surrogate AI models that provide faster approximations for design exploration and optimization.

This is more than an efficiency gain. It is a qualitative shift in the spaces engineering teams can explore. Instead of selecting the three most promising design candidates to simulate, teams can evaluate hundreds in parallel — including variations they might never have considered under traditional constraints.

Days → Min

Complex simulation time acceleration with AI Physics

1000s

Design iterations evaluated in the time of traditional single runs

3X

Expert productivity boost through automated engineering workflows

Unifying Product Data: Connected, Automated and AI Ready

Integrated, intelligent modeling and simulation environments are a material leap forward for engineering teams. But to drive AI transformation, these platforms need to become unified data fabrics—connecting data from enterprise systems and warehouses while automating data pipelines to provide an AI ready data foundation.

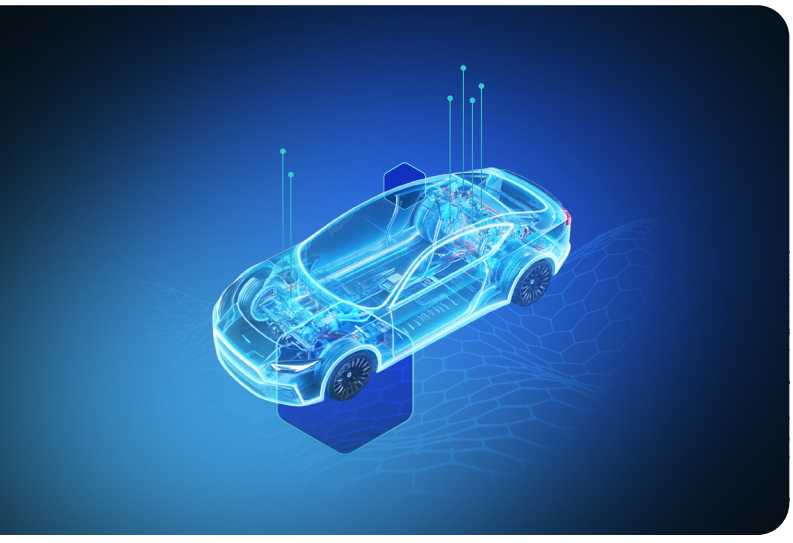

The Data Fragmentation Dilemma

The modern R&D organization generates extraordinary volumes of data: simulation outputs, CAD files, test results, maintenance records, material properties, supplier data, regulatory documents, and decades of institutional project knowledge. This data lives across siloed systems and most organizations lack the tools to effectively unify this data — impeding their ability to adopt AI-driven knowledge bases, agentic systems and surrogate models at scale.

As a result, the data that should be R&D organizations biggest advantage has actually become their biggest bottleneck to

innovation. Engineers spend an enormous amount of time finding, connecting, adding context, and structuring data rather than using it. AI models trained on poorly structured or incomplete data produce unreliable results. Intelligence that could accelerate new product development is never surfaced or leveraged.

And lack of AI-ready data is creating additional barriers. Recent BCG research[2] identifies data architecture as one of the five critical levers for AI-powered R&D success — specifically, new architectures with dedicated AI and data layers that improves data ingestion from multiple sources and facilitates rapid, sustainable AI deployments.

Unifying Product Data and Automating Insights

Intelligent data management addresses data fragmentation, context and structure issues at their root, creating a unified data fabric that connects data from PLMs, CAE environments, simulation archives, and operational systems. It automates the process of contextualizing and structuring it for use across AI driven R&D workflows so organizations can shift from data-dependencies to data-driven operations.

Key capabilities include:

- Automated metadata tagging and context enrichment as part of simulation workflows, eliminating the manual effort of data classification and improving discoverability.
- Seamless integration with existing data fabrics or with individual PLM and CAE platforms, databases, and repositories ensures data is unified as it is generated.
- Engineering knowledge graphs that capture relationships between inputs, simulations, designs, and outcomes to preserve institutional knowledge and make it queryable by both humans and AI agents.
- AI-powered data quality management that identifies inconsistencies, flags gaps, and continuously improves the reliability of the data estate.

> **"By continuously compounding and optimising our data, our intelligence and our engineering philosophies at unimaginable speed, we can deliver product developments at pace, while protecting the DNA of our company."**
>
> **Nick Collins**
> CEO, McLaren Automotive

From Data Repository to Engineering Intelligence

The transformation that intelligent data management enables is not merely operational. It is strategic. When every simulation, every design decision, and every engineering outcome is captured, contextualized, and fed back into the system, the platform becomes a continuously learning repository of engineering intelligence — one that grows smarter with every project.

Rescale builds engineering knowledge graphs from historical project data, powering agentic engineering workflows and accelerating product development decisions. Each new McLaren vehicle program, for example, builds on the accumulated intelligence of every previous one — compounding the value of decades of engineering expertise into an always-available, always-improving AI foundation.

Digital Engineering AI: Generative, Agentic and Surrogate Models

Agentic AI and AI surrogate models are not incremental improvements to existing R&D processes—they are force multipliers that redefine what is possible for every engineer.

Embedding AI: Assistants and Agent-Accelerated Innovation

Mature digital engineering platforms are moving beyond general, bolt-on LLM integrations to create embedded, domain-specific AI assistants that understand your data landscape, specific engineering processes and functional goals to provide highly tailored assistance directly in line with engineering workflows.

And Agentic AI for engineering is enabling domain-specific agents capable of semi- and fully-autonomous execution of complex, multi-step engineering tasks and workflows while maintaining observability and engineer-in-the-loop control. These are not simple automation scripts: they are reasoning systems that understand context, interpret instructions, manage tools, learn specialized skills and adapt to changing requirements.

In R&D and product engineering, agents will begin to handle more of the tasks that have historically consumed large fractions of highly skilled engineers' time:

- Automated pre- and post-processing of simulation data with quality-checking inputs, structuring outputs, and flagging anomalies before they propagate through the development workflow.
- Autonomous design space exploration that generates, evaluates, and iterates on design candidates across defined parameter spaces without continuous human direction.
- Work order generation and engineering documentation — producing technical specifications, review documents, and design reports from structured data sources.
- Integration and systems engineering support, including resolving cross-discipline data conflicts, identifying integration risks, and supporting complex analysis tasks.

BCG research identifies AI agents as a game changer for R&D activities, with documented results including 80% lead-time reductions in ship deck design, 30-40% acceleration in automotive design cycles, and up to 50% reductions in physical test costs across industries.[2]

AI Physics: Expanding Design Spaces with Surrogate Models

AI Physics represents surrogate models trained on traditional simulation data that can quickly provide approximations for the outcomes of complex physical processes at a fraction of the time and cost of full simulations. Rather than solving the underlying partial differential equations from scratch, surrogate models learn the relationship between inputs and outputs from a training corpus of an organization's own high-fidelity simulations, then provide near-instant predictions for new design configurations.

The implications for engineering productivity are profound:

- Simulation times that previously required hours or days can be compressed to seconds or minutes, enabling real-time design exploration and instant feedback on design changes.
- Surrogate models continuously improve as they are exposed to more of your simulation data, becoming increasingly accurate and reliable over time.
- Multi-physics optimization — simultaneously optimizing for aerodynamics, structural integrity, thermal performance, and other domains — becomes manageable at scale.
- R&D leadership is able to empower the entire value chain from designers to suppliers by making advanced methods broadly accessible and actionable.

McLaren is transforming its product development by embedding true end-to-end agentic AI across its entire engineering lifecycle. McLaren is leveraging Rescale's Digital Engineering Platform and NVIDIA AI infrastructure to bring unprecedented speed and scale to their design and engineering process.

This environment connects McLaren's computer aided engineering (CAE), systems engineering, and design functions into a single, continuously learning platform that adheres to McLaren's exacting quality standards and performance characteristics, while accelerating velocity.

Agentic workflows can be generated through conversational prompting, dramatically lowering the barrier to adoption and putting the power of AI-driven engineering directly in the hands of domain experts without requiring deep technical configuration. Resulting in:

- Rapid Design Exploration: Engineers can now evaluate thousands of design iterations in hours, covering multiple physics and engineering domains, accelerating optimal performing design.
- Real-Time Performance Prediction: Machine learning models enable near-instant predictions of manufacturing performance in the production of McLaren's high-performance carbon fibre structures and components.
- Agentic Engineering Embedded In Production: McLaren deployed engineering agents to perform studies scoped to specific budgets, explore hardware-software optimizations that improve cost efficiency and automate complex, repetitive tasks.

"Our foundational platform allows McLaren to leverage breakthrough agentic engineering technologies powered by NVIDIA AI infrastructure, providing a compounding source of competitive advantage for engineers in critical areas such as carbon materials, structural dynamics, and durability, and ultimately the programmatic scaling of engineering excellence across every discipline to deliver world-class products faster."

— **Joris Poort,** Founder and CEO, Rescale.

Real-world Spotlight

McLaren Automotive — one of the world's most technically demanding engineering organizations — is demonstrating what AI-Powered Engineering looks like in production.

> “AI isn’t just for efficiency anymore. It can double the pace of R&D to unlock up to half a trillion dollars in value annually.”
>
> **McKinsey,**
> *The Next Innovation Revolution—Powered by AI*, 2025

The Combined Effect: An AI-Powered Innovation Resurgence

The combination of cloud HPC with unified, intelligent data powering agentic engineering and AI physics creates something qualitatively new: an AI-first R&D operating model where the pace of discovery and development is no longer limited by human capacity or resource availability.

Engineers become orchestrators of intelligent systems rather than executors of repetitive tasks. A single engineer can supervise workloads that previously required entire teams — evaluating thousands of design candidates, running continuous simulations, and receiving AI-synthesized recommendations that surface the most promising directions for human review and refinement.

The AI-First Engineer: Redefining Engineering and R&D Operations

AI-Powered Engineering does not replace engineers — it elevates them. The highest-performing R&D organizations of the next decade will be hybrid teams where human creativity and AI-driven execution are inseparable.

Redefining the Role of the Engineer

The transformation underway in R&D mirrors the broader shift described across sectors: as AI agents take on the repetitive, data-intensive, and procedural elements of work, engineers are freed to focus on the activities that require judgment, creativity, and domain intuition.

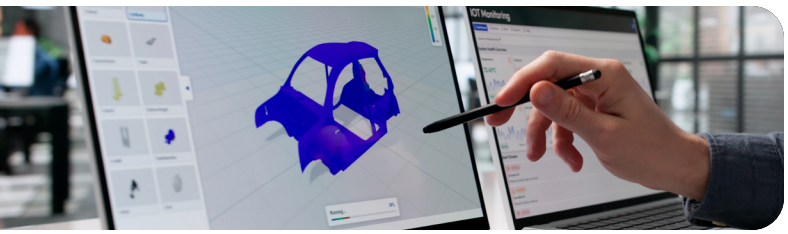

In the AI-Powered Engineering model:

- Engineers set design requirements, performance targets, and optimization objectives — defining what success looks like and why.
- AI agents execute: running simulations, processing data, exploring design space, and surfacing recommendations.
- Engineers evaluate, override, and refine: bringing domain expertise, contextual judgment, and creative insight to bear on AI-generated outputs.
- Agents learn from every human decision — improving their models, their workflows, and their understanding of what constitutes good engineering in this specific context.

The Knowledge Compounding Effect

One of the most strategically important ways AI will reshape R&D is through evergreen knowledge. In traditional R&D, institutional knowledge walks out the door with retiring engineers, resides in inaccessible archives, or is simply never captured in the first place. Each new program starts from a relatively low knowledge baseline.

In an AI-Powered Engineering environment, every project enriches a shared, searchable knowledge graph at the organizational level. Every simulation outcome, every design decision, every engineering trade-off is captured, contextualized, and always available. The result: organization gets smarter with every product development cycle and innovation accelerates as the knowledge base grows.

The Path to AI-Powered Engineering and Peak R&D Performance

AI-Powered Engineering is not a single destination. It is a continuously advancing capability that compounds in value.

The Innovation Dividend

McKinsey's research identifies AI's capacity to generate not merely more design candidates, but more diverse and unconventional ones — ideas that human engineers would be unlikely to produce given the biases of their training and experience. In the same way that AlphaGo's famous 'Move 37' was 'a move no human would ever make' and a move that proved foundational to winning, AI-powered engineering can surface similar breakthroughs.

For R&D organizations, this means a genuine expansion in the innovation frontier: not just doing what they already do faster, but discovering things they could never have discovered before.

The Journey Forward

Perhaps most importantly: AI-Powered Engineering creates continuous, compounding value. Every project enriches the knowledge graph. Every simulation improves the surrogate models. Every engineering decision teaches the AI agents to perform better.

Organizations that begin this journey earlier build a larger, more accurate, more differentiated engineering intelligence layer than those who start later. In industries defined by product differentiation — automotive, aerospace, defense, energy, and life sciences — your organization's engineering intelligence is a source of competitive advantage.

AI-Powered Engineering is how that advantage is built and sustained.

As with most journeys, two things are usually true:

1. First, not every company will follow the same path. Mature organizations with a high degree of AI-readiness may accelerate directly into scaled AI deployments, while others may take a more sequential approach.
2. And secondly, selecting a partner with both the technical and the domain expertise to support your transformation is critical.

At Rescale we work with leading engineering organizations adopting AI-Powered Engineering through three interconnected steps:

- **Establish a Flexible, Scalable Foundation:** Implementing a foundation that provides your engineering teams with the flexible, scalable computing and modern, automated modeling and simulation environments is critical for AI-powered engineering.
    - Scalable High-Performance Computing: Legacy infrastructure limits simulation frequency, fidelity, and scale. Cloud-native HPC eliminates the queue and delivers continuous, on-demand access to the latest computing architectures.
    - Integrated Simulation Environments: Fragmented tools and manual workflows create inconsistency and delays. A unified, automated modeling and simulation environment gives every engineering team access to the full simulation stack.
- **Ensure Data Is Connected And AI Ready:** Simulation and product data accumulates across disconnected systems, stripped of context and invisible to AI. Connected and context-rich data is essential for accelerating AI initiatives.
    - Unified Data and Engineering Intelligence: A unified data fabric connects simulation archives, PLM platforms, and operational systems, capturing context automatically and transforming accumulated engineering history into a continuously improving intelligence layer.
- **Elevate R&D Capabilities with AI:** Disparate AI POCs and one-off model development create excitement but typically fail to operationalize at scale. Deploying AI applications that are native to the workflows engineers already use become force multipliers.
    - Engineering Specific Agent Frameworks: High-value engineers spend the majority of their time on manual, repetitive tasks rather than the creative and analytical work. Deploying engineering agents directly into existing workflows automates basic tasks and multi-step workflows, allowing every engineer to operate at a higher level.
    - Operating System for AI Physics: Full-fidelity simulations are too slow and costly to use in tight design loops or large-scale design exploration. Surrogate models trained on your own simulation data deliver near-instant performance predictions, enabling real-time design iteration and dramatically expanding the range of ideas evaluated.

The cumulative effect of orchestrating cloud HPC, advanced modeling and simulation, data intelligence, AI physics and agentic engineering through a single, integrated platform is R&D operating at a materially higher level of performance — not just efficiently, but the volume and quality of innovative products delivered.

Organizations that build across these dimensions create an engineering intelligence layer that grows more capable with every product cycle, turning the journey itself into a compounding competitive advantage.

Sources

1. mckinsey.com/capabilities/quantumblack/our-insights/the-next-innovation-revolution-powered-by-ai

2. bcg.com/assets/2025/executive-perspectives-ai-powered-r-and-d-14feb.pdf

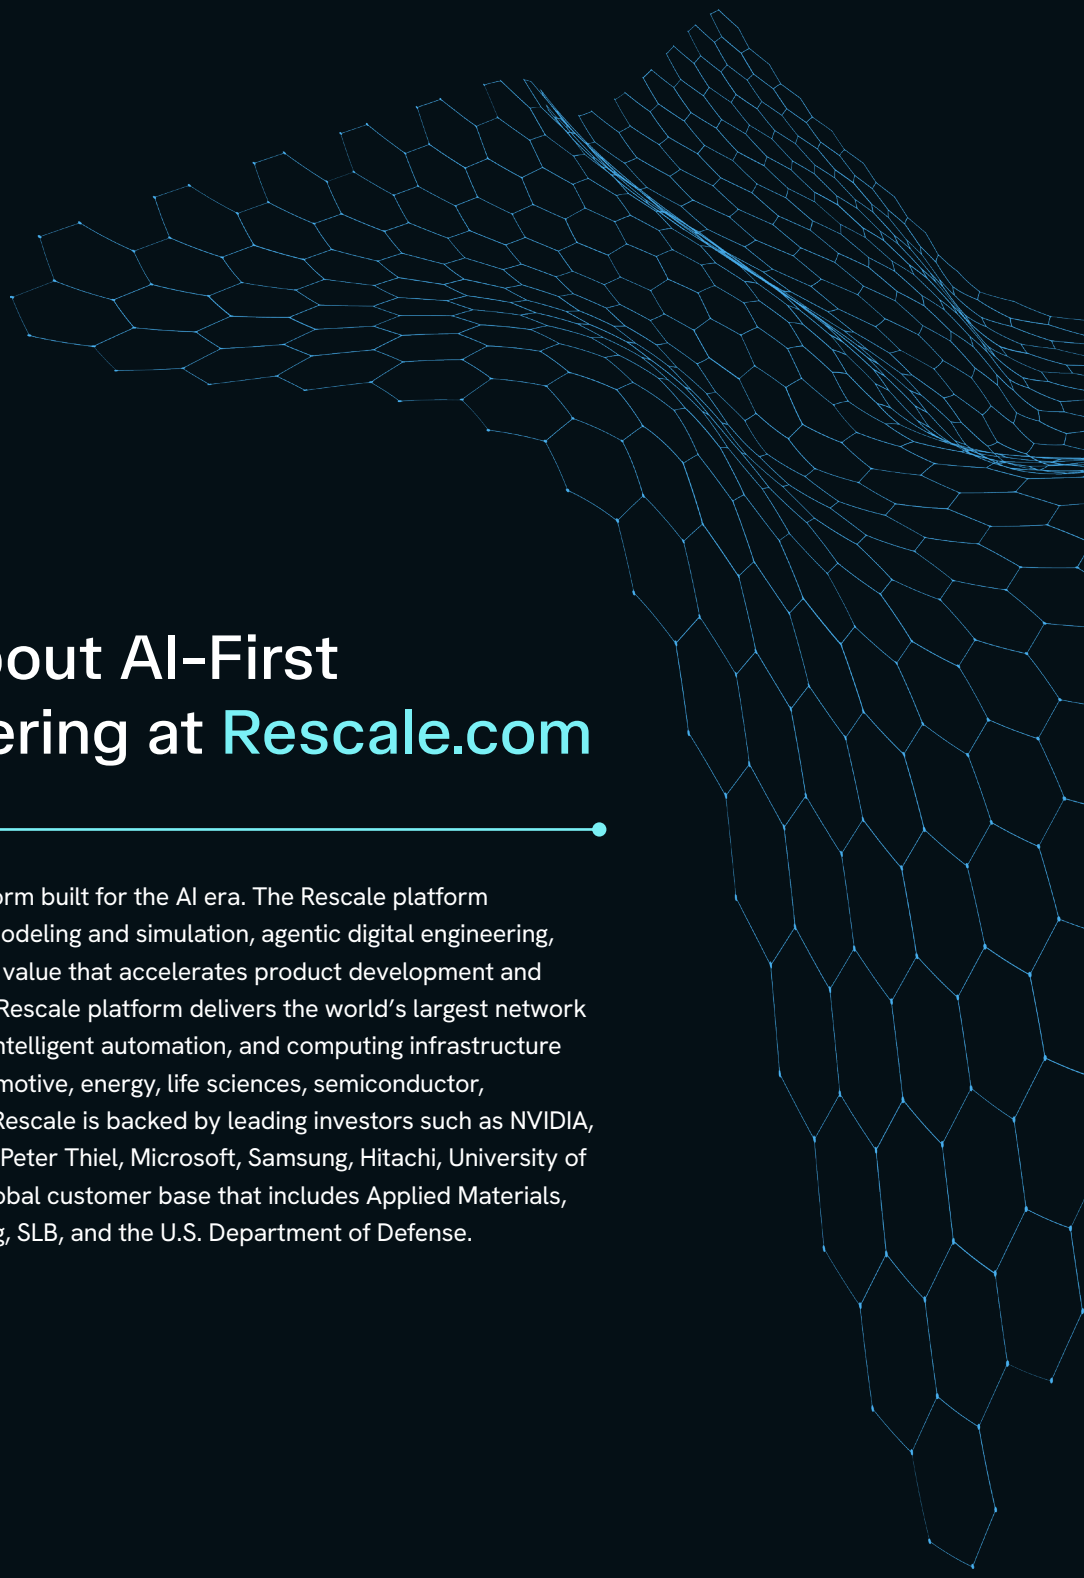

Learn More About AI-First Digital Engineering at Rescale.com

Rescale is the digital engineering platform built for the AI era. The Rescale platform integrates intelligent HPC, advanced modeling and simulation, agentic digital engineering, and AI physics to create compounding value that accelerates product development and empowers digital transformation. The Rescale platform delivers the world's largest network of engineering and R&D applications, intelligent automation, and computing infrastructure to enterprises across aerospace, automotive, energy, life sciences, semiconductor, manufacturing, and the public sector. Rescale is backed by leading investors such as NVIDIA, Sam Altman, Jeff Bezos, Paul Graham, Peter Thiel, Microsoft, Samsung, Hitachi, University of Michigan, and others. Rescale has a global customer base that includes Applied Materials, General Motors Motorsports, Samsung, SLB, and the U.S. Department of Defense.