The Prototype Problem

How AI and digital twins are changing the economics of product development

A point of view for CIOs, CTOs, CDOs and Chief Digital & Innovation Officers on where product learning should happen, what decides whether it scales, and how to put the case to a board.

THE THESIS IN ONE CHART · THE COST OF THE NEXT DESIGN ATTEMPT

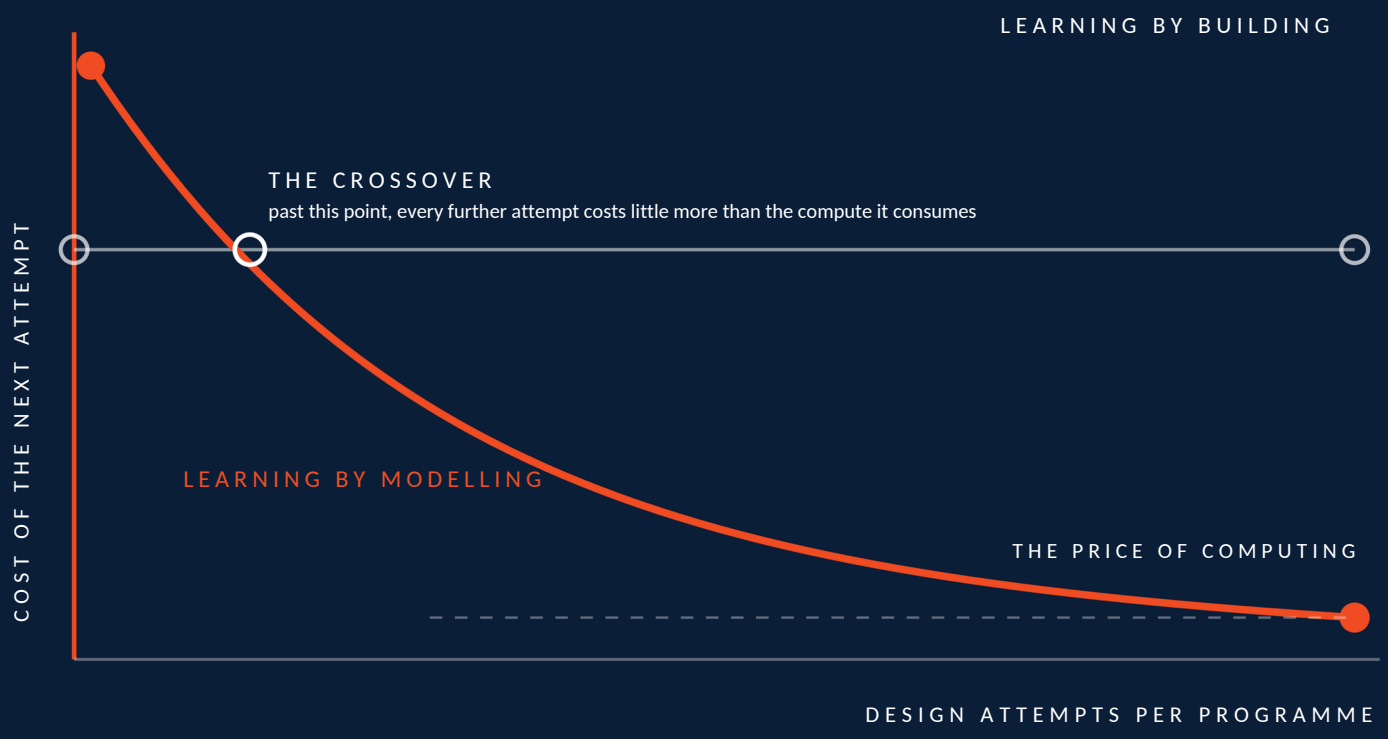

THE PAPER IN FIVE NUMBERS

80–85% of testing now done in software on Stellantis's newest platforms

50% faster to market where digital twins are used well

70% less software development time from an architecture-level twin

8% of engineering teams use surrogate models extensively; 92% have them

0.64% of GDP:India'sspend on research and development

Sources: dSPACE and Stellantis, January 2025; McKinsey & Company, April 2022; L&T Technology Services and Marelli, January 2024; SimScale, March 2026; Economic Survey 2025–26. Full references on page 11.

IN THIS PAPER

The prototype is becoming the last check, not the main teacher

Leading programmes now do most of their testing in software, so the next attempt at a design costs little more than the computing it runs on. Order and scope decide the return, not budget and the limit is readiness, not appetite.

CONTENTS

01 Why this matters now

Threepressureshaveconverged on product development inside a single product cycle. Separately, each is familiar. Together they have made the physical prototype an indefensible place to do most of a company's learning.

THE SHIFT
Complexity has moved into software and electronics

THE SQUEEZE
Launch dates have tightened while capital has not loosened

THE UNLOCK
Trying a design is no longer rationed by cost

Three pressures, one product cycle

Each of these is familiar on its own. It is their arrival together, inside a single cycle, that changes where a company can defensibly do its learning.

PRESSURE 01

Complexity has moved into software and electronics

A vehicle, an infusion pump or an industrial drivetrain is now an integration problem at least as much as a mechanical one. The number of interactions that have to be checked grows far faster than the number of parts — and interactions are exactly what a physical prototype checks worst, and latest, once architecture, tooling and supplier commitments are already sunk.

PRESSURE 02

Launch dates have tightened while capital has not loosened

Engineering leaders are asked to shorten cycles and spend less in the same planning round. Historically one was bought with the other: speed came from building more prototypes, and prototypes cost money.

Any lever that breaks that trade-off is strategic rather than merely operational.

PRESSURE 03

Trying a design is no longer rationed

Computing has become fast and cheap enough to search a range of designs rather than test a single hypothesis, and to approximate physics quickly enough that iteration is no longer limited by solver time.

That is the quiet change underneath everything else in this paper.

THE QUESTION HAS CHANGED

For decades, leadership asked: how fast can we build a prototype?

The better question is now: how much of the product can we prove before we build one?

WHY IT LANDS ON THE CIO

The first two pressures are engineering's to absorb. The third is bought, built and run as enterprise technology, which is why the ceiling on this shift is set outside the engineering function.

Nothing has changed about the difficulty of building a car, a pump or a drive train. What has changed is where the engineering team is allowed to fail.

Where you are allowed to fail

The clearest illustration is automotive. On Stellantis's newest platforms, 80–85% of testing is now done in software rather than on hardware. Engineers begin validating up to a year before hardware exists, and new vehicle software reaches customers in under two years.

80–85%
of testing now done in software rather than on hardware

12 months
of validation before hardware exists

<2 years
from start to software in the customer's vehicle

WHEN FAILURE SURFACES, AND WHAT IT COSTS

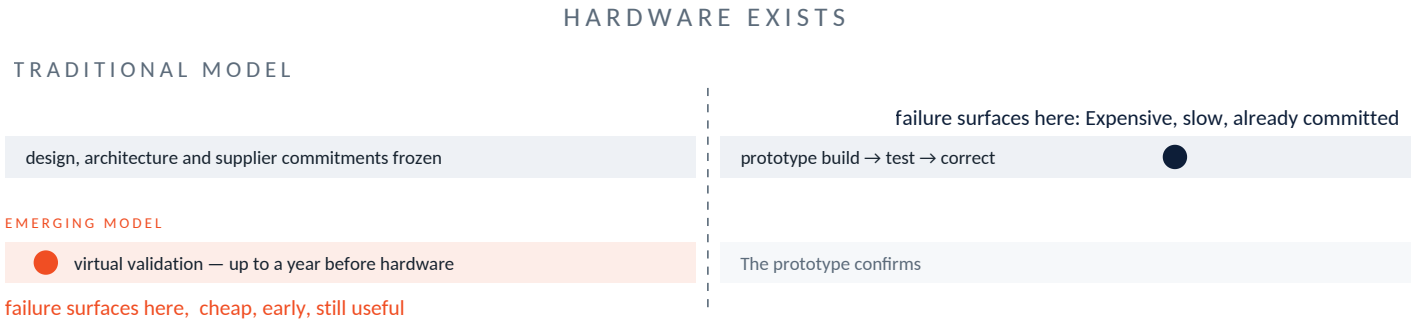

The prototype stops teaching

In the old model, failure shows up on aphysical prototype: expensive, slow, and late enough that the design is already committed. In the new one it shows up in a model, cheap, early, and still useful. The prototype stops being the thing that teaches and becomes the thing that confirms.

THE KYNDRYL LENS

Moving the point of failure earlier is a decision about sequence, not about spend, but it only holds if the estate underneath can carry continuous validation. **A programme that reorders the work without securing the data, computing and integration beneath it moves the failure earlier and the cost nowhere.**

For decades, leadership asked how fast a prototype could be built. The better question is how much of the product can be proved before one exists.

Source: dSPACE and Stellantis, Virtual Engineering Workbench customer reference, January 2025.

02 Proving the product before you build it

Traditional development treatsthephysicalprototypeasthegatebetween concept and production. The emerging model changes the order of the work — and with it, the shape of the cost curve.

THE REORDER
Design → simulate → optimise → prove virtually → prove physically

THE ECONOMICS
A change in cost behaviour, not only in speed

THE DISTINCTION
Three capabilities that business cases routinely blur

The order of the work changes first

Software now generates and compares configurations against the things that actually matter weight, strength, heat, energy use, cost to make, material used and then tests them across thousands of operating conditions without a build for each.

01	02	03	04	05	06
Design	**Simulate**	**Optimise**	**Prove virtually**	**Prove physically**	**Manufacture**
Intent, constraints and target economics set together	Behaviour checked across thousands of operating conditions	Configurations compared against competing constraints at once	The design is released against evidence from a model	The prototype confirms rather than teaches	Fewer late changes after commitment

WHERE LEARNING NOW HAPPENS

50%
sooner to market, at best, where digital twins are used well

25%
improvement in product quality, at best

WHERE CONFIRMATION HAPPENS

20–30%
of development cost absorbed by testing and validation alone

WHAT THE GAIN ACTUALLY IS
The gain is not one prototype avoided. It is every attempt that no longer needs building which is why the return does not cap out at a line item in the testing budget.

READ THE ESTIMATE HONESTLY
These are upper bounds from programmes that had the foundations in place. Treat them as evidence of what the model makes possible, not as a forecast for any single programme.

Advantage shifts from how much you can afford to test, to how quickly you can learn for what you spend.

Sources: McKinsey & Company, "Digital twins: The art of the possible in product development and beyond", April 2022; McKinsey & Company, "Testing and validation: From hardware focus to full virtualization?", 2017.

A change in cost behaviour, not just in speed

This distinction decides how the investment case is written. Building prototypes is close to linear: each attempt costs roughly what the last one did, so learning is rationed by budget. Working virtually is expensive up front, good models, clean data pipelines, serious computing, but past a certain number of iterations, each further attempt costs little more than the computing it consumes.

	LEARNING BY BUILDING	LEARNING BY MODELLING
Cost of the next attempt	Broadly the same each time	Falls towards the price of computing
Where the money goes	Recurring: builds, tooling, rigs, test time	Up front: models, data, computing capacity
What limits learning	Budget and build lead time	Model quality, data quality and trust in the result
Where advantage comes from	Size of the engineering organisation	Learning per unit of R&D spend
How it typically fails	Late design change, after commitment	Permanent pilot, bought, never relied upon

WHAT THIS MEANS FOR THE BUSINESS CASE

A case written against prototypes avoided is a cost-reduction case, and it is capped. A case written against cost behaviour is a capital-allocation case: the same budget buys a different number of proven design decisions every year that follows.

WHERE THE FAILURE MODE MOVES

Note the bottom row. Building fails late and visibly. Modelling fails quietly, the capability is bought, the licences are renewed, and no release decision ever rests on it. That failure does not appear in a technology review.

THE KYNDRYL LENS

The up-front half of this table is an enterprise technology balance sheet, not a software purchase: governed product data, engineering-grade models, hybrid cloud and high-performance computing, and the security and observability to run them at production scale. The cost curve only bends where that half has been built. Where it has not, the spend is incurred and the curve stays flat.

Building rations learning by budget. Modelling rations it by trust and trust is engineered.

What is genuinely new

Simulation is thirty years old. Its combination with modern data and modelling is not. Older tools answered *what happens if we change this?* Newer ones answer *what should we change, and which of these designs is even worth testing?* Three capabilities sit behind that, and business cases routinely blur them.

Shapes and configurations no one would have drawn by hand

Proposed against several competing constraintsat once

weight against strength, heat against cost to make, rather than optimised for one and checked against the rest.

Fast approximations of physics

Trained onhigh-fidelityruns andknownas surrogate models accurate enough to make iteration near-continuous. McKinsey's R&D Leaders Forum reports these halving development time, and reducing detail design time for steel structures from weeks

to hours.

Twins that stay connected

Linked to thefactory and the field, so the model improves with every unit made and used and so the version engineering releases against is the version manufacturing and service are working to.

EVIDENCE · MIT TECHNOLOGY REVIEW INSIGHTS, MARCH 2026

A study of 300 product, engineering and technology leaders at companies above US$500m in revenue, across 16 industries, found these tools adopted in layers, each with its own threshold of trust, with verification and getting it

right first time valued above speed.

JUDGEMENT STILL DECIDES

That is not caution for its own sake. In safety-critical work, a model no one has verified is a liability, not an asset.

Sources: McKinsey & Company, R&D Leaders Forum, 2025; MIT Technology Review Insights with L&T Technology Services, "Pragmatic by design", March 2026.

The question is no longer what happens if we change this. It is which of these designs is even worth testing.

03 Companies that changed the order of the work

Value arrives through two distinct moves: doing the proving earlier, and modelling the system rather than the part. Most programmes make one and miss the other.

MOVE ONE
Prove earlier, Stellantis, ahead of hardware

MOVE TWO
Model the architecture, not the component, Marelli

THE TEST
Whether someone outside the company will accept the evidence

CASE 01 · AUTOMOTIVE OEM · CLOUD VALIDATION PLATFORM

Stellantis: validating before the hardware exists

Stellantis · VirtualEngineering Workbench,withdSPACE
Customer reference · January 2025

THE CHALLENGE
Vehicle programmes had become software programmes, but proving still waited on hardware. Defects surfaced only once physical prototypes existed — the point at which architecture was frozen, suppliers were committed and every correction cost the most. Global teams could not test a system that did not yet physically exist, so integration risk built up quietly until late.

THE RESPONSE
Stellantis built the Virtual Engineering Workbench with dSPACE, a cloud platform on which teams around the world develop, integrate, test and tune vehicle software before hardware is available, running continuously rather than in gated cycles.

THE OUTCOME

80–85%
of testing now runs in software on new platforms

12 months
engineering startsup toa year before hardware is available

100×
faster development cycles credited to the Virtual Cockpit

New vehicle software reaches customers in under two years.

WHAT TRANSFERS
The return came from moving the proving earlier, not from buying more simulation. Reordering the work is the least expensive decision available here, and the one most often skipped.

THE KYNDRYL LENS
Continuous validation is a continuous workload. A platform that runs in gated cycles can be sized for peaks; one that runs continuously has to be architected for availability, cost and security together — which is where the ambition usually meets the estate.

Source: dSPACE and Stellantis, Virtual Engineering Workbench customer reference, January 2025.

CASE 02 · TIER-ONE SUPPLIER · SYSTEM-LEVEL TWIN

Marelli: modelling the architecture, not the component

Marelli · Cloud digital twinofthevehicleE/Earchitecture,withL&TTechnologyServices
Press release · January 2024

THE CHALLENGE
Marelli, a global tier-one supplier, faced rising software content spread across cluster, infotainment, zone control units and body domains. Component-level testing could confirm that each piece worked, but not that the vehicle's electrical and electronic architecture behaved correctly as a whole. Proving that required physical prototype vehicles — which made both timelines and prototype spend very hard to compress.

THE RESPONSE
With L&T Technology Services, Marelli built a cloud-based digital twin of the entire electrical-electronic architecture rather than of individual components.

THE OUTCOME

70%
less software development time

30%
lower prototype costs

WHAT TRANSFERS
Scope sets the size of the prize. A twin of a part gives local optimisation; a twin of the architecture changes how the product gets developed. Worth noting who built it, Indian engineering capability delivering measured outcomes for a global tier-one supplier.

THE TWO MOVES, SIDE BY SIDE
Stellantis moved the when. Proving shifted ahead of hardware.

Marelli moved the what. The unit of modelling shifted from component to architecture. Neither required a larger testing budget.

THE KYNDRYL LENS
An architecture-level twin is an integration problem before it is a modelling problem: it needs design, software, test and supplier data reconciled into one operating picture. Scope is set by the data that can be brought together, not by the ambition of the model.

Source: L&T Technology Services and Marelli, "LTTS and Marelli develop cloud-based digital twin of vehicle electrical-electronic architecture", January 2024.

Beyond automotive

The mechanism travels.In two industries with nothing else in common, the constraint was never ambition it was the cost of an attempt.

CONSUMER GOODS

Hindustan Unilever · Doom Dooma, Tinsukia, Assam

WEF End-to-End Value Chain Lighthouse · January 2025

An AI-enabled digital twin lets teams simulate and test sustainable packaging options before committing them to the line, turning a trial from a line event into a computation.

- 2→30 — trialsperyear,2019 to 2023
- −84%
- −21% — virginplasticuse at the site

PHARMACEUTICALS

Dr Reddy's · Bachupally, Hyderabad

WEF End-to-End Value Chain Lighthouse · 40+ use cases

Digital twins deployed alongside advanced analytics and industrial IoT across more than 40 use cases. The percentages matter less than where they were earned an industry in which every material change has to survive a regulator.

- 43% — improvement in cost of manufacturing
- −30% — productionlead time

THE THIRD CONDITION · ADMISSIBILITY

A third condition decides whether any of this counts: whether someone outside the company will accept the evidence. In November 2023 the US Food and Drug Administration issued final guidance on how it assesses the credibility of computational modelling in medical device submissions, aligned to the ASME V&V 40 standard.[7] In regulated industries the obstacle was never whether the model was accurate. It was whether the model was *admissible*. Treat that conversation as part of the business case, not a compliance step at the end the same holds in aerospace, rail, energy and anywhere a certifier, customer or insurer has to sign.

In regulated industries the obstacle was never accuracy. It was admissibility and that is a governance question, settled early or not at all.

Sources: World Economic Forum and Unilever, January 2025; Dr. Reddy's Laboratories, October 2022; US Food and Drug Administration, final guidance, November 2023.

04 Why most programmes stop at the pilot

The marked as bought the capability.Very few depend on it. The explanation is not solver quality, it is what the solvers are being asked to run on.

THE GAP	THE EXPLANATION	THE ANSWER
92% have surrogate models; 8% use them extensively	Infrastructure readiness and governance maturity, not architecture	Five conditions that decide whether this scales or stalls

Owned, not relied upon

A survey of 350 senior engineering leaders at companies with more than 1,000 employees across the US, UK and Germany, published by SimScale in March 2026, gives the clearest picture of where the market actually sits.[8]

THE ADOPTION GAP · SIMSCALE, MARCH 2026

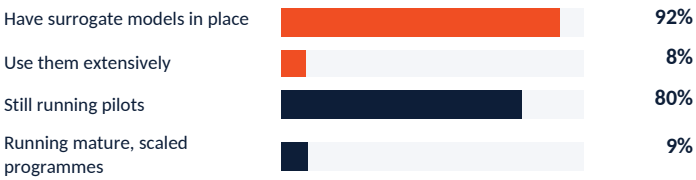

Survey of 350 senior engineering leaders at organisations with 1,000+ employees across the US, UK and Germany.

- 3–4× — as many design options evaluated per programme
- 6 hrs — simulation turnaround, against seventeen

THE FINDING THAT MATTERS

The important finding is not the gap but its explanation. Whether a company scaled tracked more closely with the readiness of its infrastructure and the maturity of its governance than with its data architecture.

The organisations that scaled did not have better solvers. They had somewhere solid to run them.

HOW TO READ THIS IN YOUR OWN PORTFOLIO

Count the release decisions in the last twelve months that rested on a model rather than a rig. If the answer is none, the capability is owned and not relied upon, whatever the licence count says.

The gap is not between companies that have these tools and companies that do not. It is between owning something and depending on it.

Source: SimScale, "State of Engineering AI 2026", March 2026.

Five conditions that decide whether this scales

A twin is only as good as the data, model sand infrastructure beneath it. Skipping one does not slow the programme it caps it.

WHAT IT TAKES	WHY PROGRAMMES FAIL WITHOUT IT
Product data you can trust	Without governed design, simulation, test, bill-of-material and field data, models learn from fragments and twins decay into pictures nobody will stake a release on.
Models built for engineering	Safety-critical design needs models that work alongside physics-based simulation and established validation — not general-purpose office tools pointed at engineering problems.
Computing that scales	Hybrid cloud, high-performance and GPU capacity for heavy simulation and training — with speed, security and cost managed as one architecture rather than three separate purchases.
A connected thread to the factory	Engineering cannot work in one version of the product while manufacturing, quality, procurement and service hold different ones.
Rules people believe in	Model validation, data lineage, IP protection, explainability, cybersecurity and human sign-off — decisive wherever a regulator, customer or insurer must accept evidence that came from a model.

THE KYNDRYL LENS

None of the five is a feature of a simulation package; each is a property of the estate the simulation runs on. The product may be designed by engineering, but the ability to design it this way is an enterprise technology problem, which is the sentence that decides who sponsors the investment.

NONE OF THIS IS A SHOPPING LIST

A gap in any one of the five is not a delay to be worked around. It is a ceiling and it should be priced into the business case before approval, not discovered after it.

THE CHEAPEST QUESTION TO ASK FIRST

Not “which capability do we want next?” but “which single condition is capping us today, and what is the cheapest path to removing it?”

The organisations that scaled did not have better solvers. They had somewhere solid to run them.

05 India: more learning per rupee

Researchcapitalisscarce.Engineeringcapabilityisnot.That combination makes this a capital-efficiency question in India before it is a technology question.

THE CONSTRAINT

0.64% of GDP spent on R&D; business funds only 41% of it

THE STAGE

The “valley of death” between research and market-ready product

THE REFRAME

Proven design decisions per unit of capital, not cost per engineering hour

Scarce research capital, abundant capability

The Economic Survey 2025–26 places India’s innovation gap exactly where this paper begins, the “valley of death” between early research and market-ready products, where prototyping, piloting and characterisation stall.[9] That is precisely the stage this approach compresses.

GROSS EXPENDITURE ON R&D, % OF GDP

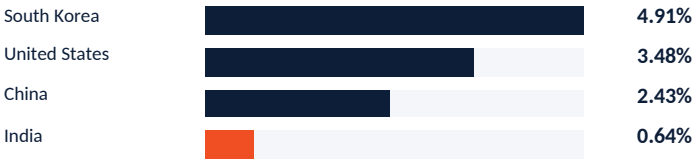

Business contributes only 41% of India’s national R&D total. Source: Economic Survey 2025–26.

THE CAPABILITY ALREADY IN PLACE

~650
models Tata Steel runs across its global value chain, with three sites in the WEF Lighthouse Network

US$63b
engineering R&D services within a US$315b technology sector in FY26

17% → 22%
India’s share of global engineering R&D sourcing, FY24 to 2030

When research capital is scarce but engineering talent is not, the question sharpens: how do we make every rupee generate more learning, faster? Advantage will belong to the fastest learners, not the largest R&D organisations.

WHAT THIS REFRAMES

For Indian companies and for the capability centres and engineering service providers running global programmes from India this reframes the pitch. The claim is no longer cost per engineering hour. It is proven design decisions per unit of capital, and that claim has to be demonstrable rather than asserted.

Advantage will belong to the fastest learners, not the largest R&D organisations.

Sources: Economic Survey 2025–26; Financial Express Futech, December 2025; Nasscom Strategic Review 2026; Nasscom and Everest Group, 2024.

06 The Kyndryl view: the ground a twin stands on

Sections 04 and 05 arrive at the same place from opposite directions. The five conditions that separate scaled programmes from permanent pilots are not features of a simulation package. They are properties of the technology estate the simulation runs on.

THE QUESTION FOR THE CIO
Can the estate carry a workload the product roadmap now depends on?

THE VANTAGE POINT
Mission-critical estates run in more than 60 countries

THE PATTERN
The modelling is rarely what holds a programme back

Buying it is no longer the hard part; relying on it is

Kyndryl's vantage point is the layer beneath the engineering platform, mission-critical technology estates run for customers in more than 60 countries. Seen from there, the pattern is consistent. The modelling is rarely what holds a programme back. What sits underneath it usually is.

EVIDENCE · KYNDRYL 2026 PEOPLE READINESS REPORT, 1,100 LEADERS ACROSS EIGHT COUNTRIES

57%
running these tools broadly or embedded in core processes, up from 35% a year earlier

32%
had met even one of their top two objectives

11%
had met both

In engineering, the same shape appears: 92% of teams have surrogate models, 8% lean on them.[8] Two separate datasets, one conclusion,the gap is between owning something and depending on it.

Virtual proof is a question of trust before it is a question of computing

81%
expect autonomous systems to make consequential business decisions within twelve months

25%
fully trust such systems without human oversight

33%
have clear rules on which decisions those systems may take

THE KYNDRYL LENS

Releasing a design on the strength of a model is one of those decisions. Where a programme cannot say which model is authoritative for which decision, at what confidence, and who signs it off, engineers go back to hardware — and the economics go back with them.

Trust is not a cultural problem downstream of the technology. It is the mechanism by which the cost curve is realised.

Sources: Kyndryl, "2026 People Readiness Report", 25 June 2026; SimScale, "State of Engineering AI 2026", March 2026.

The foundation is older than the ambition sitting on it

Heavy simulation, model training and keeping a twin in step with the real thing are among the most demanding workloads a company will ever run. They find weakness in the estate before anything else does.

EVIDENCE · KYNDRYL 2025–2026 SECURITY AND NETWORKS SNAPSHOT, 3,700 LEADERS ACROSS 21 COUNTRIES

25%
of mission-critical networks, storage and servers are at end of service

37%
of leaders judge their network ready for what is coming

35%
are investing heavily in upgrades[14]

WHAT THIS MEANS IN PRACTICE

The foundation has to be engineered around the outcome — connecting simulation and modelling work to the data, hybrid cloud, cybersecurity, observability and integration needed to run it reliably, repeatedly and at production scale.

IN THE FIELD · DETROIT

Kyndryl's collaboration with Wayne State University, announced on 13 January 2026, established the IntelliMake research hub and pilot-scale factory — combining advanced computing, cybersecurity and digital twins, with Kyndryl Bridge orchestrating the infrastructure behind the line.[15]

IN MANUFACTURING ENGAGEMENTS

Twin and predictive-maintenance foundations built on hybrid data fabric and edge have delivered:[16]

+15%
machine availability

−10%
maintenance spend

WHERE KYNDRYL FITS

The role is not to replace the design and simulation platforms companies already own. It is to build and run the backbone around them — so that the cost of the next attempt genuinely does fall towards the price of computing, and stays there across product generations.

A quarter of mission-critical infrastructure is already at end of service. The most demanding workload a company will ever run is the one being planned on top of it.

Sources: Kyndryl, "2025–2026 Security and Networks Snapshot", 10 March 2026; Wayne State University and Kyndryl, 13 January 2026; Kyndryl manufacturing client outcomes.

07 The case to put to the board

The cases could not open with"howmanyprototypes can we eliminate?" That framing caps the answer at a line item. It should rest on five outcomes — and the strongest cases carry all five.

THE FRAME
Five outcomes, each converting capability into a number the board already tracks

THE DISCIPLINE
A scorecard, not a summary: baseline, target, review

THE TEST
Evidence all five and it is a business case; evidence one and it is a purchase

Five outcomes, not one line item

These five are the bridge from technology to business value, and they are the form in which the programme should be governed. Each converts an engineering capability into a number the board already tracks.

OUTCOME	THE QUESTION TO ANSWER	WHERE IT SHOWS UP
Faster to market	How many weeks come out of design, validation and change cycles?	Revenue timing; competitive position
More from the same engineers	How many more design options can the same team evaluate?	Options tested per programme
Less spent on building and testing	How much physical testing, tooling, scrap and rework is avoided?	Direct development spend
Lower product risk	How many defects are caught before physical testing, ramp-up or launch?	Warranty, recall and ramp-up cost
Better use of capital	Can spend shift from repeating validation to genuine innovation?	R&D mix; innovation capacity

USE IT AS A SCORECARD, NOT A SUMMARY
Baseline each outcome before the programme starts, commit to a target on at least three, and review them at the cadence applied to any other capital decision.

THE SEQUENCING IS INSTRUCTIVE
Speed to market and engineering leverage arrive early and compound. Lower product risk and better use of capital appear only once models are trusted enough to release against — which is why the readiness conditions in section 04 are not a hurdle to clear before the value starts, but the thing being bought.

A programme that can evidence all five is a business case. One that can evidence only the third is a procurement request.

Four points for the C-suite

The goal is not to eliminate the physical prototype. It is to make it the final check rather than the main teacher, moving experimentation into a place where the next attempt costs close to nothing.

1 The economics change, not just the tooling

Twins can bring products to market as much as 50% sooner and lift quality by up to 25%,[2] against testing and validation costs of roughly 20–30% of the total.[3] The prize is the cost curve, not a line item.

2 Order and scope beat spend

The largest returns came from proving earlier — 80–85% of testing moved into software ahead of hardware[1] — and from modelling the architecture rather than the part, cutting software development time by up to 70%.[6]

3 Readiness is the constraint, not ambition

92% of engineering teams have surrogate models but only 8% use them extensively, and 80% are still piloting;[8] across the wider enterprise, 57% have these tools in core processes yet only 11% have met both their top two objectives.[13] The twin is not the product. The data, the models, the computing and the rules beneath it are.

4 For India, this is a capital question

At 0.64% of GDP spent on R&D,[9] advantage will belong to the fastest learners, not the largest R&D organisations.

THE CLEAREST SIGNAL OF A MATURE PROGRAMME

It is not a larger engineering software budget. It is a release decision that rested on a model, a named owner who signed it, and a CFO who can see the outcome in a number the board already tracks.

This is not an engineering innovation alone. It is a C-suite technology and business agenda — sponsored, funded and measured against the five outcomes rather than against a prototype count.

THE BOTTOM LINE

The prototype is not disappearing. It is being demoted, from the main teacher to the last check.

Every attempt that no longer needsbuilding is an attempt that costs closeto nothing. That is a change in cost behaviour, not a saving, and it compounds with every product generation that follows.

The companies that capture it will not be the ones that bought the most simulation. They will be the ones that built the ground it stands on first — trusted data, engineering-grade models, computing that scales, a connected thread to the factory, and rules people believe in.

Advantage will belong to the fastest learners, not the largest R&D organisations, and the speed at which an organisation learns is now a property of its technology estate.

THREE QUESTIONS TO LEAVE THE ROOM WITH

01

How much of our next product could we prove before we build one — and what is stopping us?

02

Which release decision in the last year rested on a model rather than a rig, and who signed it?

03

Of the five conditions, which one is capping us — and what would it cost to remove it?

Anoop Kubba
Customer Advisory Head,
Mint CIO Circle

Sources

All figures are as published by the cited sources. Access dates are August 2026.

1 dSPACE and Stellantis, Virtual Engineering Workbench — customer reference on Stellantis's cloud-based software development, integration and test platform, January 2025. Development, integration, testing and optimisation up to a year before hardware availability; 80–85% of testing on software-in-the-loop platforms on new technology platforms; vehicle software delivery in under two years; Virtual Cockpit accelerating development cycles by up to 100 times versus traditional methods.

2 Mickael Brossard, Sebastien Chaigne, Jacomo Corbo, Bernhard Mühlreiter and Jan Paul Stein, "Digital twins: The art of the possible in product development and beyond", McKinsey & Company, 28 April 2022. Time to market accelerated by as much as 50%; product quality improved by up to 25%.

3 McKinsey & Company, "Testing and validation: From hardware focus to full virtualization?", 2017. Testing and validation typically account for 20–30% of product-development cost.

4 McKinsey & Company, R&D Leaders Forum, 2025. Deep-learning surrogate models trained on high-fidelity simulation data enabling rapid iteration and halving development time; generative design reducing detail design time for steel structures from weeks to hours.

5 MIT Technology Review Insights, in partnership with L&T Technology Services, "Pragmatic by design: Engineering AI for the real world", March 2026. Survey of 300 product development, engineering and technology leaders at organisations with revenue above US$500m across 16 industries.

6 L&T Technology Services and Marelli, "LTTS and Marelli develop cloud-based digital twin of vehicle electrical-electronic architecture", press release, January 2024. Software development time reduced by up to 70%; prototype costs reduced by up to 30%.

7 US Food and Drug Administration, "Assessing the Credibility of Computational Modeling and Simulation in Medical Device Submissions", final guidance, November 2023. Risk-informed credibility framework aligned to ASME V&V 40.

8 SimScale, "State of Engineering AI 2026", March 2026. Survey of 350 senior engineering leaders at organisations with 1,000+ employees across the US, UK and Germany; approximately 3–4× more design variants per programme; simulation turnaround of six hours versus seventeen; 92% deploying surrogate models with 8% extensively; 80% running pilots against 9% with mature, scaled programmes.

9 Government of India, Ministry of Finance, "Economic Survey 2025–26", 2026. Gross expenditure on R&D at 0.64% of GDP (US 3.48%, China 2.43%, South Korea 4.91%); business sector contributing 41% of total; the prototyping and piloting "valley of death".

10 Financial Express Futech, "Tata Steel's AI playbook", December 2025. Approximately 650 models across the global value chain; three manufacturing sites in the World Economic Forum Lighthouse Network.

11 Nasscom, "Technology Sector in India: Strategic Review 2026", released at the Nasscom Technology & Leadership Forum, 24 February 2026. Technology sector projected at US$315b in FY26 (6.1% growth), of which engineering R&D services account for US$63b.

12 Nasscom and Everest Group, "The Global ER&D Shift", 2024. India's share of global engineering R&D sourcing rising from 17% (FY24) to 22% by 2030.

13 Kyndryl, "2026 People Readiness Report", released 25 June 2026. Global study of 1,100 senior business and technology leaders across eight countries and six priority industries; fieldwork by Edelman Intelligence, 24 March to 30 April 2026. 57% deploying broadly or embedding in core processes, up from 35%; 32% achieving at least one of their top two objectives and 11% achieving both; 81% expecting autonomous agents to make impactful business decisions within twelve months; 25% fully trusting such systems without human oversight; 33% with clear policies on decision authority.

14 Kyndryl, "2025–2026 Security and Networks Snapshot" within the Kyndryl Readiness Report, published 10 March 2026. Insights from 3,700 business and technology leaders across 21 countries; 25% of mission-critical networks, storage and servers at end of service; 37% judging their network ready for future risk against 35% investing heavily in network upgrades.

15 Wayne State University and Kyndryl, "Wayne State University and Kyndryl announce collaboration to advance AI-driven manufacturing and workforce innovation in Detroit", press release, 13 January 2026. IntelliMake research hub and pilot-scale factory; advanced computing, cybersecurity and digital twins, orchestrated with Kyndryl Bridge.

16 Kyndryl, manufacturing industry client outcomes (internal engagement data). Digital twins, predictive maintenance and hybrid data fabric/edge delivering a 15% increase in machine availability and a 10% reduction in maintenance expenditure.

17 World Economic Forum, "Global Lighthouse Network 2025", press release, 14 January 2025; and Unilever, "Five ways Unilever's new Lighthouse site applies AI", January 2025. Hindustan Unilever's Doom Dooma site, Tinsukia, Assam; AI-enabled digital twin for sustainable packaging trials; trials from two per year in 2019 to 30 in 2023; time per trial reduced by 84%; virgin plastic use reduced by 21%.

18 Dr. Reddy's Laboratories, "World Economic Forum recognises Dr. Reddy's Hyderabad factory as part of its Global Lighthouse Network", press release, 11 October 2022. Bachupally, Hyderabad site; more than 40 Fourth Industrial Revolution use cases across six technologies including digital twins, advanced analytics and industrial IoT; 43% improvement in cost of manufacturing; 30% reduction in production lead time.