

# Rethinking the operating model for product, IT, data and AI

---

*Architecture, governance and execution in the agentic era.*

*Across the organisations we work with, product, IT, data and AI no longer share a single view of what the enterprise is actually betting on. Each function runs its own perimeter through its own scattered tools. The operating model that would tie them together does not, in any practical sense, exist as a system: it is reconstructed by hand at every steering committee, drawn from sources that quietly diverge. That missing system is the subject of this paper.*

# Reading note

---

*This paper maps the terrain on which Chief Data, Analytics and AI Officers, Chief Product Officers, and Heads of Portfolio increasingly need to operate. The diagnosis is built from three converging sources: what McKinsey, BCG and Gartner have published over the past eighteen months; what the European regulator has now codified; and what we observe in our own work with portfolio and platform teams.*

The EU AI Act was designed to bring high-risk systems under a defined set of obligations on 2 August 2026. In May 2026, the Union deferred that deadline to December 2027—not because the regulatory ambition has shifted, but because neither member states nor most enterprises were ready. None of the obligations have been weakened. Only the date has moved. And it has moved because no one yet operates the layer that would make daily compliance demonstrable. That layer is precisely the subject of this paper. Part 2 returns to the deferral in detail.

The reading time is approximately forty minutes. Every figure is sourced, and the full bibliography is consolidated in Annex A: McKinsey, BCG, Gartner, PwC, the World Economic Forum, RAND, Silicon Valley Product Group and Product Talk.

## POSITION OF THIS DOCUMENT IN THE FYGURS SERIES

This founding paper establishes the thesis of convergence and the reference architecture. Three supplements will follow, addressed in turn to product managers, heads of PMO, and CDAIOs. This document should be read first; the supplements build on it explicitly.

# The argument in one page

---

Six observations, summarised below.

**1. Product, IT, data and AI are converging into a single operating model—integrated and value-driven.** McKinsey's *The Agentic Organization* (2026), BCG's *AI Radar* the same year, and Gartner's reclassification of the *Project & Portfolio Management* category as *Strategic Portfolio Management*—three independent vantage points—arrive at the same conclusion.

**2. The deferral of the EU AI Act is the signal, not the news.** The regulation was scheduled to apply to high-risk systems on 2 August 2026; in May 2026, the Union confirmed a shift to December 2027, citing the unreadiness of states and enterprises alike. The European register, conformity assessments, human oversight and data governance obligations remain intact. Only the deadline has moved. A sixteen-month deferral granted because no one was ready says more about the underlying problem than the original deadline did. Demonstrating compliance will require something other than a slide deck.

**3. The gap between adoption and value is now the story.** RAND finds that 80% of AI projects failed to deliver expected value in 2025. McKinsey reports that 88% of organisations are using AI in at least one function, while only 39% can point to a measurable impact on EBIT. The 49-point gap is the principal signal of the cycle.

**4. The diagnosis reduces to three familiar symptoms.** Feature factories, the dominance of spreadsheets, and the "three-PM test"—each is well-documented and none has been resolved by a decade of agile, SAFe or OKRs.

**5. AI governance frameworks stop at the doorstep of the product organisation.** Most 2026 frameworks describe an AI Office, a use-case lifecycle and a risk playbook with discipline. Very few describe how the work lands inside product teams. That, however, is where the value is created or lost.

**6. The architecture that works rests on four layers.** Bet, opportunity, initiative and experiment—deployed within three containers (workspace, portfolio, product) and read through three lenses (product manager, PMO, CDAIO). One entry, three readings.

# Convergence has already happened

---

*Two facts impose themselves at the outset, and neither is a forecast: both have already happened. Gartner has redrawn a market category in full, recasting Project & Portfolio Management (PPM) as Strategic Portfolio Management (SPM). At the C-suite level, a single role is consolidating—the CDAIO—pulling together responsibilities long distributed across data, IT and AI. A third signal follows: in 2026, McKinsey, BCG and Gartner publish their reference frameworks within months of one another, and when read together they describe the same operating model. This part traces the terrain.*

## McKinsey · Five pillars, one system

In *The Agentic Organization* (March 2026), McKinsey identifies five pillars: the business model, the operating model, governance, people and culture, and technology and data. The framework refuses, immediately, to read them in silos. The argument fits on a single line: these are not five compartments to coordinate; they are the facets of one system.



*Technology leaders must codevelop solutions while aligning on architecture, governance, and risk from the outset to capture the full value of the agentic era.*

MCKINSEY, *THE AGENTIC ORGANIZATION* (2026)

## BCG · The work happens outside the algorithm

BCG's *AI Radar 2026* reaches the same conclusion through a different angle: AI transformation succeeds only when it is anchored in strategy, embedded in redesigned processes, and carried through to scale.

Its 10-20-70 framework distributes the effort without ambiguity—10% to the algorithm, 20% to technology and data, 70% to people and processes. The work, in other words, does not sit inside the model. It sits inside the operating model that surrounds it.

*AI transformation must be linked to strategy, built into redesigned processes, and adopted at scale.*

CHRISTOPH SCHWEIZER, CEO BCG (AI RADAR 2026)

## Gartner · The end of PPM, the start of SPM

The most explicit signal comes from Gartner. The legacy *Project & Portfolio Management* category has been formally rebranded *Strategic Portfolio Management* (SPM)—a transition codified by the 2025 Magic Quadrant and the 2026 roadmap infographic. The PMO of the previous decade, whose first question was « are we delivering on time? », gives way to a portfolio function whose first question is now « are we placing the right bets? ». Gartner is no longer describing the same function.

The market follows. Planview, ServiceNow, Planisware and six other vendors are now evaluated on their ability to bring interdependent portfolios, strategy modelling and operational execution onto a single platform.

## Leadership · Data, analytics and AI under one authority

In parallel on the HR side, the consolidation is just as visible. The Chief Data, Analytics and AI Officer (CDAIO) is absorbing roles that were until recently fragmented: Chief Data Officer, Chief Technology Officer for data, Head of AI. The logic is straightforward: data cannot be governed without governing the AI that feeds on it, and AI cannot be governed without governing the products that embed it. The consolidation is happening at the top, and HR observers are documenting it in real time.

## Synthesis

Taken one at a time, these signals would read as competing recommendations from rival houses. Taken together, they describe the same operating model.

- **Integrated governance** (McKinsey) covering business model, operations, talent, data and AI.
- **An effort centred on people and processes** (BCG: 70% of the effort), not on the algorithm.
- **A unified strategic portfolio** (Gartner SPM) replacing legacy siloed PPM.
- **Consolidated leadership** (CDAIO) spanning data, analytics and AI, with explicit articulation to product and IT.

This is the terrain the paper documents. The next part addresses why the moment has arrived; the three that follow describe what we observe on the ground, what is missing, and what works.

# Why inaction is now expensive

---

*Three forces have made standing still costly: European regulation, the persistent gap between adoption and value, and the shift to agentic AI. None is new. What is new is their simultaneous alignment on the same calendar.*

## **The regulatory clock · The deadline has slipped by sixteen months**

Regulation is the first of these forces, and the most tangible: it puts a date on the gap this paper describes. The EU AI Act (Regulation (UE) 2024/1689) enters into force in phases. The phase that was scheduled to affect most enterprises fell on 2 August 2026: from that date, high-risk systems would have been subject to a binding set of obligations—registration in the European register, conformity assessment, risk management, data governance, logging, human oversight and transparency.

That deadline has moved. In November 2025, the Commission tabled the Digital Omnibus on AI as an amending instrument; after the failure of an initial trilogue on 28 April, a political agreement was reached in early May 2026 and confirmed by the Council on 13 May. The revised calendar is unambiguous: high-risk standalone systems (Annex III) shift to 2 December 2027, and high-risk systems embedded in regulated products (Annex I) to 2 August 2028. A sixteen-month slippage on the principal milestone.

Two clarifications are warranted. First, as of writing (June 2026), the agreement remains in formal adoption. Until it is published in the Official Journal, the original text remains the legal reference, and 2 August 2026 the binding date. Compliance teams largely plan against December 2027, but the legal switchover has not yet occurred. Second, not everything is deferred: certain transparency obligations on deployers remain due on 2 August 2026, while the marking of generated content has been pushed to December 2026.

The reason for the deferral is the most telling element of all. The Union has not stepped back on substance; it has acknowledged that implementation is not ready—harmonised standards still in draft, national authorities not all designated, enterprises behind. Not a single obligation has been softened. Only the date has slipped, and it has slipped because the ecosystem has failed to deliver, on time, the operational layer that those obligations presuppose. That is the gap this paper examines.

The fines, by contrast, are unchanged—calibrated, deliberately, to put the topic on the executive committee's agenda: up to €35M or 7% of global revenue for prohibited practices, and up to €15M or 3% for non-compliance with obligations attached to high-risk systems.

**What this changes in practice.** Sixteen extra months do not alter the nature of the work. An AI Office documented in PowerPoint will not survive scrutiny from a regulator, an internal client, or an audit committee. Each will demand evidence: a traceable validation procedure for every high-risk use case, documented production-time oversight, and a conformity assessment that can be retrieved at any moment. This is portfolio work, and the sixteen months will reward only those who use them for it.

## AI is everywhere; its value, much less so

The pattern is now well-documented: most AI investment has not delivered the value it promised. Three recent studies put numbers on the gap.

- **RAND (2025):** more than 80% of AI projects fail to deliver expected business value—roughly twice the failure rate of comparable IT projects. The report identifies five root causes, foremost among them the disconnect between business objectives and the actual capabilities of the AI.
- **McKinsey Global AI Survey (2026):** 88% of organisations now use AI in at least one function. Only 39% report a measurable EBIT impact. The 49-point gap is the principal signal.
- **MIT (2025):** 95% of GenAI pilots never reach production.

What these figures collectively say is straightforward: AI works, but the organisation around it does not keep pace. What McKinsey and BCG label "people and processes" is precisely what is missing between pilot and production.

## With agents, the leap becomes organisational

Through 2024, enterprise AI was predominantly predictive (models) and generative (assistants). Through 2025 and 2026, the centre of gravity shifts to the agentic: systems that decide, trigger actions, and operate in closed loops.

McKinsey makes it the title of its report, *The Agentic Organization*; BCG flags it as the structuring inflection of the decade. The leap is less technological than organisational. Who supervises an agent that decides? Who carries the liability? Who measures its value?

All of these are questions of operating model. None resolves at the level of the algorithm.

# The same obstacles, a decade later

---

*In our work with product, IT and data teams struggling to extract value from AI, the same three symptoms recur with striking regularity. None is novel: the product literature has documented all three for ten years. All three nevertheless dominate the landscape in 2026, because no transformation programme has yet uprooted them.*

## Symptom 1 · The three-PM test

The test is intentionally crude: in the same week, with no preparation, convene three product managers from the same organisation, ask each what the company is betting on this quarter, and compare the answers.

If they diverge, there is no strategy. There are three private interpretations of one. A strategy that no team can recite is no longer a strategy: it is a file.



*Most product orgs have a sophisticated execution layer and an undefended strategic call. That gap is the work now.*

FYGURS, WHAT MATTERS EDITION 07

### THE TEST IN ONE SENTENCE

If the three PMs were replaced tomorrow, would the new trio tell the same story? When the answer is no, the strategy lives in the heads of those leaving, not in the system that remains.

## Symptom 2 · The feature factory

The term belongs to Marty Cagan, threading through *Inspired* (2017), *Empowered* (2020) and *Transformed* (2024). A feature factory measures its success in story points shipped—never in business metrics that move.

The pattern is unmistakable at a glance: the roadmap accretes through accumulated requests; no one writes hypotheses that could be invalidated; and experimentation remains the exception. The organisation counts what it ships, never what that shipping changes.

*It's not about the number of features you build, but about the value those features deliver.*

MARTY CAGAN, SILICON VALLEY PRODUCT GROUP

On the AI side, the feature factory takes a distinctive turn: a backlog of use cases accumulates—"we have forty-seven in the queue"—without any clear hierarchy between what genuinely moves the business and what is essentially technical curiosity.

### **Symptom 3 · The dominance of Excel**

In the large majority of organisations we work with, there is an Excel file, maintained by the PMO, the CDAIO, or a senior analyst. It is the true source of truth on the portfolio, and everyone knows it: the executive committee references it, monthly reviews project it.

Yet that file has none of the properties of a system: no audit trail, no reliable history, and collaboration that consists of email attachments. It is a personal artefact, and it collapses the day its author leaves the organisation.

The dominant Excel is not, in itself, a failure—it is a patch, filling the void left by the absence of a dedicated software system. Removing it too quickly invites chaos; entrenching it institutionalises a human bottleneck. The right path is to replace it with a platform that covers the same use cases while restoring what the spreadsheet cannot offer: auditability, collaboration, and history.

### **The three symptoms reinforce one another**

Considered in isolation, each symptom is treatable. Considered together, they feed each other: in the absence of a real strategy, the product organisation ships features for lack of direction; the only remaining view of the whole sits in a spreadsheet. It is this knot that a portfolio operating system has to untie, first by surfacing the diagnosis, then by equipping a different mode of work.

# Every framework omits the same layer

---

*Most large European enterprises now operate an AI governance framework, or are deploying one through 2026. The diagrams circulate: produced by internal directorates, consulting houses, or vendors. The form varies; the structure converges.*

We have redrawn the typical framework (Figure 1)—not to reproduce it, but to separate what works from what is missing.

## The typical structure

At the top sits the AI governance layer, organised in three blocks: the governance framework (principles, policies, risks, ethics); the operating model (committees, articulation with IT and risk); and the AI Office, accountable for steering, design authority and business-side liaison.

At the centre, an AI factory organised in six phases: discover, design, develop, validate, deploy, operate. A risk playbook runs end-to-end across the cycle, on top of a shared foundation of data, tooling and infrastructure.

At the base, five or six strategic axes: alignment, risk control, industrialisation, collaboration, value measurement, continuous learning.

## What the typical framework gets right

Three credits deserve to be acknowledged.

- **Governance as a distinct layer.** Without it, EU AI Act compliance does not hold.
- **The lifecycle structured in phases.** This is the only known way to industrialise beyond the pilot stage.
- **The AI Office as a defined function,** not a virtual coordinating cell. This single distinction changes everything for transverse coherence.

Many enterprises have not yet assembled these three; getting there is already a meaningful step.

## What the typical framework misses

Three omissions stand out, in ascending order of gravity.

**The bridge to product teams.** The framework treats AI as a separate atelier. In practice, modern products blend AI and non-AI work. A use case is only valuable if it lands in a product, owned by a product team, and judged on the outcomes that product is accountable for. The bridge between the AI Office and the product organisation is rarely drawn—so it gets built case by case, on goodwill.

**What defines a bet.** The framework speaks of use cases, but never specifies what separates a good bet from noise: no falsifiability criterion, no three-PM test, no framing of expected outcomes. The risk playbook is rich. The value playbook is missing.

**The watchtower in PowerPoint.** The watchtower appears, formally, as the observability and value-tracking layer—but it lives in slideware far more often than in code. On any Tuesday morning at nine, no one knows where the forty-seven use cases in the portfolio currently stand, which are at risk, which are creating value, which are quietly drifting. The map stays silent, because the map exists only as a slide.

## How to use the typical framework

The framework is not wrong; it is incomplete. The governance layer is solid; the industrialisation layer is approximately specified; and the operational layer, where most of the actual work lives, is simply absent. That is the layer the rest of this paper documents.

# Product and AI in a single architecture

---

*Here it is. The architecture rests on four layers of objects, three containers and a single source of truth. Every piece of work described in the governance framework above orchestrates through these objects. None of the objects is invented here: each one draws on the product or AI literature. Their articulation, however, is ours.*

## The four layers

This architecture extends a strategy stack we have published previously. The three original layers—"the bet, the structure, what we ship"—become four, in order to distinguish the initiative from the experiment that validates it.

The four layers read top-down, from strategic to operational.

- **BET.** A strategic assertion that is falsifiable, measurable, and dated. A bet describes what one believes strongly enough to commit capacity to—and is willing to see invalidated. Without a bet, there is no strategy. There is a budget.
- **OPPORTUNITY.** A customer need, problem, or business lever that contributes to the bet. The concept originates with Teresa Torres's Opportunity Solution Tree (*Continuous Discovery Habits*, 2021).
- **INITIATIVE.** The solution explored to address one or more opportunities. This is what product and engineering teams carry day-to-day.
- **EXPERIMENT.** A concrete test of one hypothesis embedded in the initiative. An experiment can succeed, fail, or remain inconclusive. It is the one level at which invalidation is explicitly accepted.

## The three containers

Above the four layers, three hierarchical containers carry the organisational context.

- **Workspace.** The organisation as a whole. One workspace per enterprise.
- **Portfolio.** A coherent grouping by business unit, market, or mission.

- **Product.** A team and a scope. This is where bets sit.

Bets live at the product level—not at the portfolio level (too broad) and not at the initiative level (too narrow). This rule, easy to state, removes a great deal of ambiguity in the organisations we observe.

## Why the architecture holds against AI

The AI use case referenced in Part 4 is not a new object in this architecture: it is an initiative. It lands in a product, serves a bet, addresses an opportunity, and is validated through experiments.

This homogeneity produces two practical consequences.

- Product teams and data or AI teams operate on the same vocabulary. There is no longer a translation layer between two systems.
- The total portfolio—product initiatives and AI use cases combined—becomes visible in a single view. The executive committee arbitrates on one common grid.

# One truth, three lenses

---

*Part 5 established the structure. There remains a question: who looks at what, and with what concern in mind? Our answer fits a formula: one source of truth, three lenses. The lenses do not translate the underlying data; they illuminate the same objects from three angles.*

## The three roles

Three figures carry three operational questions.

- **The Product Manager** (« **Eva** ») leads a product team. Her daily question: « what are we learning this week? ». Her lens surfaces the Opportunity Solution Tree, the experiments under way, outcomes by opportunity, and the experimentation calendar.
- **The Head of PMO** (« **Pierre** ») orchestrates the portfolio. His question: « are we holding the trajectory we committed to? ». His lens surfaces initiatives by phase, dependencies, capacity, milestones, and the variance between trajectory and commitment.
- **The CDAIO** (« **Sarah** ») owns the data and AI strategy. Her question: « which are creating value, and which are diverging? ». Her lens surfaces the AI watchtower: use cases by phase, data quality, drift, EU AI Act compliance, and captured value.

## Why three lenses rather than one universal view

The appeal of a single universal view is strong, but it fails on contact every time. The PM finds it too broad; the PMO too superficial; the CDAIO too product-centric. The lens does the work: without it, the data is correct but unusable.

Take a single bet. From the PM lens, the hypotheses to test are what surface. From the PMO lens, the resources to mobilise. From the CDAIO lens, the model risk and the value captured.

## The cardinal rule · No data captured twice

If a team has to re-enter the same information to populate another lens, the architecture has failed. Entry is unique; the lens calculates. It is this seemingly modest rule—the guarantee

of uniqueness—that finally displaces the dominant Excel, far more decisively than any single additional feature could.

# The watchtower in production, not in slideware

---

*The watchtower is the most ambitious promise of the typical AI governance framework, and it is also the promise that most often remains stranded in slideware. This part describes what a functional watchtower looks like: a system in motion, not a metaphor.*

## Four observability layers

An operational AI watchtower covers four layers, ordered from the most immediate to the most strategic.

- **1. Lifecycle.** Where does each use case stand in the six phases? How many are blocked? For how long? This is the first layer executive committees ask to see.
- **2. Compliance.** Which use cases fall under "high risk" within the EU AI Act? Is the conformity assessment up to date? Is human oversight documented? This is the layer that protects the enterprise at the high-risk deadline (December 2027 under the revised calendar), and immediately for transparency obligations that already apply.
- **3. Technical quality.** Input data quality, model drift, performance decay, operational alerts. This is the layer that data and MLOps teams know well—MLflow, Evidently, Weights & Biases.
- **4. Value.** What business value is each use case actually capturing relative to the bet it serves? This is the most difficult layer, because it forces a connection between a business metric and a technical use case. And it is also, precisely, the layer that justifies the investment.

## What distinguishes a real watchtower from a slide

Three operational criteria do the work.

**REAL-TIME ACCESSIBILITY**

The executive committee, the CDAIO, or an external auditor can open the screen at any time and read the state of the portfolio—no export, no slide to reconstruct, no waiting. If restituting that state takes a day of PMO work, it is not a watchtower.

**ACTIONABLE GRANULARITY**

Clicking on a use case opens its lifecycle, dependencies, open risks, owner, and captured value. A living object that answers questions—not a static slide.

**CONTEXTUAL ALERTS**

The system surfaces, on its own, every use case beginning to drift: phase stalled too long, model drift, expired conformity, declining value capture. The alert reaches the right owner with the right context. No one chases information anymore.

**Watchtower and MLOps · Two distinct functions**

The confusion is frequent, though the two are complementary. MLOps covers technical quality—the third layer above—necessary, but insufficient for what we describe here. The watchtower, by contrast, aggregates all four layers and, crucially, reconnects them to bets. It is the bridge between technical quality (on the MLOps side) and strategy (on the bet side).

Without that bridge, observability remains confined to data teams—invisible to the executive committee and to the regulator alike.

# No one rebuilds an operating model in a quarter

---

*Whether we are looking at product transformations as documented by Cagan, or AI transformations as documented by BCG, experience converges toward a three-stage trajectory. Here is what we observe when it succeeds.*

## Months 1 to 3 · The pilot bet

Choose a bet, a team, a product. One. Not three, not ten. Preferably in a domain where AI carries real weight; without that, the three-lens test does not fully play its role.

Deploy the full architecture across that restricted perimeter: the bet, its opportunities, initiatives and experiments, viewed through all three lenses. The objective is simple—demonstrate that the system answers Eva's, Pierre's and Sarah's three questions without any parallel spreadsheet surviving.

### PHASE 1 SUCCESS CRITERIA

The team stops maintaining the parallel Excel on the pilot bet. The PMO can reconstitute the state to the executive committee in under ten minutes without preparation. The CDAIO can demonstrate the use case's compliance under the EU AI Act.

## Months 4 to 6 · The portfolio

Extend the approach to every bet in a single portfolio. This is the second wave of organisational work: harmonising vocabulary, retiring the parallel spreadsheets, redesigning the rhythms—monthly executive review, quarterly portfolio review. It is the most delicate phase, because it is the one in which legacy habits resist most strongly.

### PHASE 2 SUCCESS CRITERIA

Every bet in the portfolio is captured in the system. The executive committee accepts the watchtower as the official source. Monthly reviews happen inside the tool, not on reconstructed slides.

## Months 7 to 12 · The full operating model

Finally, extend the approach to every portfolio. This is the point at which the economies of scale materialise: cross-portfolio comparisons, executive arbitration on consistent grids, a consolidated AI watchtower, systematic EU AI Act compliance.

It is also the point at which the culture turns. New joiners absorb the vocabulary from the outset, rather than as a layer added on top of an older way of working.

### PHASE 3 SUCCESS CRITERIA

A new PM or a new CDAIO can take up the role and understand the state of the portfolio in a day—not three months. The three-PM test (Part 3) returns the same answer for each trio.

## What causes the transition to fail

Three traps recur.

- **Trying to transform everything at once.** The pilot bet is not a compromise; it is a precondition. Without it, phase 2 begins in doubt and loses momentum.
- **Confusing tool and operating model.** Software does not transform an organisation; it accelerates a transformation that has already been decided. Installing the system without altering the rituals is the surest way to add one more spreadsheet.
- **Leaving the AI Office isolated.** If the AI Office is not wired into the product organisation by phase 1, phase 3 will reveal a gap that will be expensive to close.

## What we are building

This document describes a category of software that does not yet have a stable name. We call it a portfolio operating system. Not another PPM, nor a layer of AI governance grafted onto the product organisation. It is a single source of truth on bets, opportunities, initiatives and experiments, read through three lenses—PM, PMO, CDAIO—and anchored to a watchtower that ties value back to bets. The whole serves an integrated, value-driven operating model.

This is what we are building at Fygurs. This document does not demonstrate the platform; that is not its purpose. It sets out the terrain and a reference architecture, both of which anyone is free to take and use. Confronting your own operating model against this framework—locating where it holds and where it lacks a system—is the kind of conversation we are interested in. [saad@fygurs.com](mailto:saad@fygurs.com).

# Sources

---

Every quantitative or structural claim in this document is traceable to one of the sources below. Direct links to original publications are provided where available.

- 1. McKinsey · The Agentic Organization: Contours of the Next Paradigm for the AI Era**  
2026. [mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization](https://mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-agentic-organization)
- 2. McKinsey · The State of AI 2026 (Global AI Survey)**  
November 2025. [mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai](https://mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai)
- 3. BCG · AI Radar 2026: As AI Investments Surge, CEOs Take the Lead**  
2026. [bcg.com/publications/2026/as-ai-investments-surge-ceos-take-the-lead](https://bcg.com/publications/2026/as-ai-investments-surge-ceos-take-the-lead)
- 4. Gartner · Magic Quadrant for Strategic Portfolio Management (formerly PPM)**  
2025–2026. [gartner.com/en/documents/6807234](https://gartner.com/en/documents/6807234)
- 5. Gartner · Infographic 2026 SPM Frameworks, Processes and Tools Investments Roadmap**  
2026. [gartner.com/en/documents/7393830](https://gartner.com/en/documents/7393830)
- 6. RAND Corporation · Why AI Projects Fail**  
2025. [rand.org/pubs/presentations/PTA2680-1.html](https://rand.org/pubs/presentations/PTA2680-1.html)
- 7. European Parliament and Council · Regulation (EU) 2024/1689 ("EU AI Act")**  
Original calendar: 2 February 2025, 2 August 2025, 2 August 2026, 2 August 2027. Revised by the Digital Omnibus on AI (political agreement May 2026, formal adoption ongoing): Annex III high-risk deferred to 2 December 2027, Annex I to 2 August 2028. [ai-act-service-desk.ec.europa.eu](https://ai-act-service-desk.ec.europa.eu)
- 8. Marty Cagan · Transformed: Moving to the Product Operating Model (SVPG)**  
Wiley, 2024. See also *Inspired* (2017) and *Empowered* (2020).
- 9. Teresa Torres · Continuous Discovery Habits**  
Product Talk LLC, 2021. [producttalk.org/opportunity-solution-trees/](https://producttalk.org/opportunity-solution-trees/)
- 10. Fygurs · What Matters Edition 07, "What is your team actually betting on?" (Saad Amrani Joutey)**  
May 2026. Fygurs note. The original source of the three-PM test and the strategy stack (the bet, the structure, what we ship), both extended in this paper. [fygurs.com](https://fygurs.com)
- 11. Capgemini · AI Act in Focus: AI Governance Framework for Safe Deployment**  
2025. [capgemini.com/insights/expert-perspectives/ai-act-in-focus-ai-governance-framework](https://capgemini.com/insights/expert-perspectives/ai-act-in-focus-ai-governance-framework)
- 12. Sopra Steria Next · Blueprint for Scaling Generative AI**  
2025–2026.

**13. MIT · State of AI in Business 2025 (95% of GenAI pilots fail to reach production)**

2025.

**14. PwC · 2026 Digital Trends in Operations Survey**

2026. [pwc.com/us/en/services/consulting/supply-chain-operations/library/digital-trends-operations-survey.html](https://www.pwc.com/us/en/services/consulting/supply-chain-operations/library/digital-trends-operations-survey.html)

**15. World Economic Forum · Why Data Readiness is Now a Strategic Imperative for Businesses**

January 2026. [weforum.org/stories/2026/01/why-data-readiness-is-now-a-strategic-imperative-for-businesses](https://www.weforum.org/stories/2026/01/why-data-readiness-is-now-a-strategic-imperative-for-businesses)

**16. Microsoft · Establish an AI Center of Excellence (Cloud Adoption Framework)**

2026. [learn.microsoft.com/en-us/azure/cloud-adoption-framework/ai/center-of-excellence](https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/ai/center-of-excellence)