DOCUMENT 54 · ORGANISING THE COMPANY FOR AI

How a company organises itself for AI

Who does what in real organisations, where AI governance should sit, where use cases come from, how they are governed through to production and how the information reaches the board; what organisation can be anticipated once AI is embedded and how the people of an AI-first company are organised

Fernando García Varela Version 0.8 (working draft) 19-11-2026 Draft for review

6	4	8	5
functions someone has to perform	organisation patterns observed	sources use cases come from	organisation options to become AI-first

IN THIS DOCUMENT

01 Purpose and scope

02 The six functions someone has to perform

03 What is usually seen in the market

04 From the real organisation chart to the SEVEN-G bodies

05 Where use cases come from

06 How what has entered is governed

07 How it is monitored and how the information reaches the board

08 What organisation comes next: the author's hypothesis

09 The human organisation of an AI-first company

10 Questions for a useful conversation

11 How it is set up in practice

12 Associated tools and templates

13 Related documents

14 Version control

Legal notice and disclaimer. SEVEN-G is a reference methodological framework provided "as is" and for information purposes only. It does not constitute legal, regulatory, financial or professional advice, nor does it guarantee compliance with any law or standard. References to general regulation (such as the EU AI Act, the GDPR, DORA or NIS2), to technical standards and to regulation specific to each sector or jurisdiction may be incomplete, may not apply to a particular case or may become out of date as a result of regulatory changes, interpretations or supervisory positions after their consultation date. **Each organisation that uses SEVEN-G is solely responsible for identifying the regulation that applies to it, verifying that it is current and certifying its own regulatory compliance**, with the appropriate qualified advice. This methodology is a generic, free aid shared with the community so that nobody has to start from scratch; each person or organisation can and should adapt it to its own use. It must not be inferred that its legally sensitive parts have been reviewed by legal counsel: those reviews, for each company or sector, are the ultimate responsibility of the company, consultant or organisation that uses it. Although every effort is made to keep it up to date, some regulation may have changed without being reflected here. To the fullest extent permitted by law, the author accepts no responsibility whatsoever for the effects of its application in any organisation or for its full applicability. The methodology does not grant certification of any kind. The author accepts no liability for any use made of this content or for any decisions taken on the basis of it. The data, figures, companies and cases in the examples are fictitious or illustrative.

THE ESSENTIALS Reference

Explanatory document. It places the framework inside a company's real organisation chart: which functions someone has to perform, which roles usually take them on today, where AI governance is best placed, where use cases come from and how they are governed and reported up to the board. Section 8 is forward-looking and is marked as the author's hypothesis; section 9 proposes how to organise the people of an AI-first company. It creates no rules or evidence: if it differs from documents 01, 14, 30 or 60, those prevail.

Obligation matrix and layered reading (document 94)

01

Purpose and scope

1.1 What this document answers

When a company decides to use AI seriously, the same conversation always appears: *who owns this?* IT has its own director, security another, data another; there is a transformation function, sometimes an AI function, and the board asks for information that nobody is quite sure who should prepare. This document puts that conversation in order through six questions:

1. Which **functions** does someone have to perform, whatever they are called in the organisation chart? (section 2)
2. What is seen **in the market today**, and which organisation patterns work and which fail? (section 3)
3. Where should **AI governance** sit: with the AI function, with transformation, with risk, with the board? (section 3.4)
4. Where do **use cases** come from, how are they governed and how is the **board informed**? (sections 5, 6 and 7)
5. What organisation comes next, once AI is embedded and **functional silos** stop being the natural way of dividing work? (section 8)
6. What does being an **AI-first** company mean, and how are its people organised without creating one more silo? (section 9)

1.2 What it is not

- **It is not a mandatory organisation chart.** SEVEN-G does not require creating any new function. It requires the functions in section 2 to have a name and an owner, and those that must be separated to be separated (document 01, section 8.2; document 30, section 5).
- **It does not replace document 30**, which defines the bodies, their mandates, the delegation matrix and escalation. Here we explain how those bodies fit into an organisation chart that already exists.
- **It is not a market study.** What section 3 describes are patterns observed in the author's professional practice, expressed qualitatively. It carries no percentages or market shares because they would not come from a verifiable source (document 04, section 6).
- **Section 8 is a reasoned hypothesis**, not a forecast. It is marked as such and accompanied by the signals that would allow one to check whether it is happening.
- **Section 9 is the author's proposal**, based on checked external sources. It is not a market study and does not repeat third-party savings figures.

> **Why it matters.** Most of the blockages attributed to technology are in fact questions of allocation: two areas believe they hold the same remit, or nobody holds it. A governance framework that is not translated into the real organisation chart remains a document.

02

The six functions someone has to perform

Regardless of size, sector and existing roles, AI requires **six functions** to have an owner. They may be concentrated in few people (in a small company, one person may perform three), but none can be left empty and two pairs cannot fall to the same owner.

The six AI functions in a company

Who sets direction, who proposes, who coordinates, who builds, who controls and who verifies

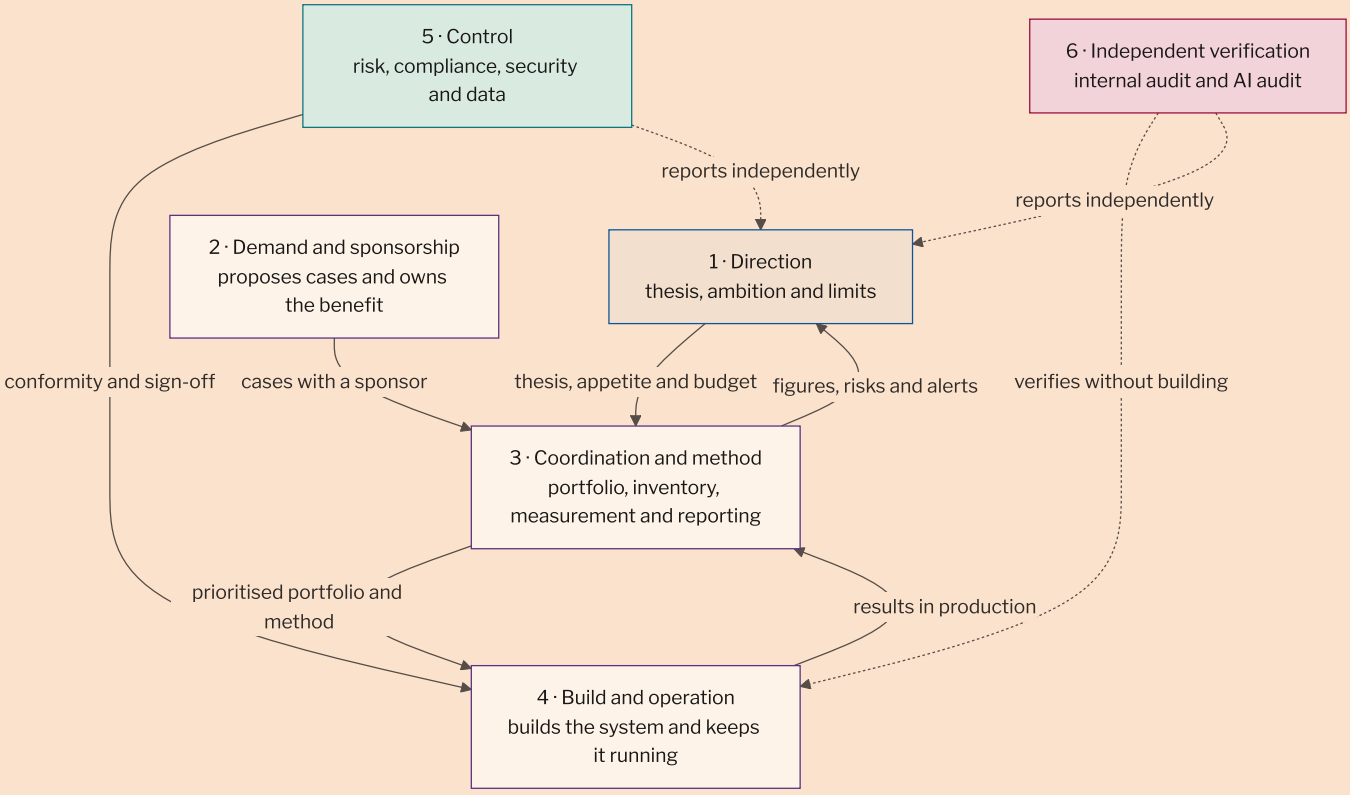

Source: SEVEN-G

#	Function	What it decides or signs	Who usually performs it	Where it sits in SEVEN-G
1	**Direction**	AI thesis, ambition by sphere, risk appetite, framework budget, approval of Transform bets.	Board and senior management.	document 13; document 30, section 3.1
2	**Demand and sponsorship**	Which problem is solved, who owns the benefit and who takes on the change in their area.	Business and operations directors.	document 01, section 8; document 43
3	**Coordination and method**	Portfolio, inventory, application of the method, consistent measurement, preparation of reporting to the bodies.	AI Office (or the function that acts as one).	document 30, section 3.4; document 14
4	**Build and operation**	Architecture, data, development, suppliers, go-live and ongoing operation.	Technology, data and the areas that operate the system.	document 52; document 53; document 51
5	**Control**	Risk assessment, regulatory conformity, security, data protection and Enterprise go-live sign-off.	Risk, compliance, security and data protection (second line).	document 30, section 9; document 33; document 35
6	**Independent verification**	Verification that the framework is applied and that the evidence is real.	Internal audit and AI Auditors (third line).	document 38

The three separations that are not negotiable (document 01, section 8.2, and document 30, section 5):

- whoever **builds** does not verify their own evidence;
- whoever **proposes and sponsors** does not sign the risk or compliance conformity;
- whoever **audits** does not take part in the build or in the *gate* decision.

> **Why it matters.** Almost every organisation that "cannot get past pilots" has functions 1, 2 and 4 well covered and functions 3, 5 and 6 without an owner. Without coordination there is no portfolio; without control there is no defensible go-live; without verification, value figures do not survive a question from the board.

03

What is usually seen in the market

This section describes **observed patterns** from practice. It is not a study with a sample or a statistic: it is a qualitative ordering, useful for placing oneself and for each company to recognise its own case. Where regulatory obligations are cited, the official source and document 34 are referenced.

3.1 Which part of AI each existing role takes on

Common role	Which part of AI it usually takes on	Where it usually falls short
IT or technology function (CIO, CTO)	Platforms, models, integration, technical security, suppliers, infrastructure cost. It is the one that can build fastest.	The value hypothesis and adoption: it delivers systems that work and nobody uses; it measures activity, not benefit.
Security function (CISO, CSO)	Model and data security, identities, access, attacks with and against AI, incidents.	It comes in late, once the solution has been chosen; with no prior criteria it can only veto.
Data function (CDO)	Quality, lineage, permitted use, data governance, sometimes advanced analytics.	It confuses data governance with AI governance: they are different and complementary (document 51).
Transformation or digital function	Portfolio of use cases, change management, relationship with the business, benefit measurement.	It rarely has authority over technology or control; it depends on persuasion.
AI function (CAIO), where it exists	AI strategy, AI Office, use cases, relationship with model providers.	If it also builds and also controls, it breaks the separation of duties; with no budget, it is an evangelisation role.
Finance function (CFO)	Budget, investment, return measurement, management control, recurring cost.	It applies project return criteria to transformation bets and blocks them (document 14, section 4.4).
Legal, compliance and data protection	Regulatory classification, impact assessments, supplier contracts, transparency.	It arrives at the last phase; without early classification, redesign is expensive (document 32).
People function	AI literacy, effect on work, employee relations, use of AI in decisions about people.	It treats AI as training rather than as a redesign of work (document 50).
Operations and business units	The problem, the process, the real data, adoption and the benefit.	They propose solutions already chosen ("we want a copilot") instead of measurable problems.
Internal audit	Verification of the framework, annual plans, follow-up of recommendations.	Without published criteria it audits perceptions; it needs the framework to exist first (document 38).

3.2 Four organisation patterns

Pattern	How to recognise it	It works when	It fails when
A · Technology leads	AI is a programme of the IT function; the business contributes cases and validates.	Technical maturity is low and the foundations have to be built: platform, data, security.	Success is measured in models deployed; the benefit has no owner and adoption does not happen.
B · Transformation leads	AI is one line of the transformation agenda, with a portfolio and change management.	There is a cross-company portfolio and the main problem is prioritising and making adoption happen.	There is no technical or control authority: the portfolio advances on paper and stalls at production.
C · AI Office, with or without an AI function	A central function exists for method, inventory, measurement and support; the business sponsors and technology builds.	There are enough use cases for the method to pay off and the central function has budget and access to the governing bodies.	The office also builds and also controls: it loses its independence and becomes a bottleneck.
D · Federated by default	Each area buys and experiments on its own; there are dozens of uses and no inventory.	Never in a stable way; it serves to discover real demand for a few months.	The first incident, the first audit or the first consumption invoice arrives: nobody knows how many systems exist or which data they touch.

In practice, most companies recognise themselves in a mixture: an official pattern and pattern D as the underlying reality. The first task in implementation is usually to **make visible what already exists** (census and inventory, document 32 and document 14, section 11) before deciding on structure.

3.3 Is an AI function needed?

Four questions help decide this better than the debate about whether the role is fashionable:

1. Are there **more use cases than a committee can follow** once a month without losing the detail?
2. Is there **more than one unit** competing for the same technical capacity or the same budget?
3. Does the company have, or will it have, **high-risk systems, agents with autonomy or direct exposure** to customers?
4. Can anyone answer within a week **how many AI systems exist, which data they use and how much they cost**?

With two or more uncomfortable answers, the central function (pattern C) pays for itself. Whether it is called AI function, AI Office or AI coordination matters less than three conditions:

- **its own budget** for discovery and method, however small;
- **direct access** to the management committee and to the board body that oversees AI;

- **not being the one who builds** the systems whose evidence it later verifies.

An AI function fails, almost always, in one of three ways: it becomes a pilot factory with no benefit owner; it becomes an approval committee with no capacity to execute; or it accumulates functions 3, 4 and 5 of section 2 and leaves the company without independent control.

3.4 Where AI governance should sit

This is the most frequent question and it has an answer of principle: **governance cannot report to whoever executes**. Function 5 (control) and function 6 (verification) report to the board or its delegated committee, not to the owner of the portfolio. That said, there are three reasonable locations and one that is not.

Option	What it consists of	Advantage	Risk	Condition for it to work
1 · Governance in the second line (risk or compliance), with technical secretariat in the AI Office	Policy, appetite and conformity sit with risk; method and inventory sit with the AI Office.	Clear independence; fits the three lines model and board oversight.	Distance from the business: control arrives late unless it is embedded in the phases.	Risk takes part from phase 0, not at the final gate.
2 · Governance in the AI function, with conformity kept separate	The AI function owns method, portfolio and inventory; risk and compliance sign-off stays with the second line.	A single point of contact for the business; speed.	Concentration: if it also builds, separation is lost.	The AI function does not build the systems or, if it does, does not verify their evidence.
3 · AI Committee as the governing body, with no dedicated function	A committee with management, technology, risk, data and people decides; the secretariat is held by whoever coordinates.	Inexpensive and sufficient in mid-sized companies; shares responsibility.	Without a dedicated secretariat, the committee neither prepares nor follows up anything.	Someone has real dedicated time to prepare decisions and follow up agreements.
4 · Governance inside the function that builds	The same area proposes, builds, approves and declares the value.	None that survives a review.	Judge and party: the evidence is not defensible before audit or before the board.	It is not an acceptable option in SEVEN-G.

In the first three options the board keeps what it cannot delegate: approving the thesis and the risk appetite, approving Transform bets and overseeing results (document 30, section 3.1). Regulation reinforces this idea in sectors where the management body already has an express responsibility for technology and resilience risks, such as Regulation (EU) 2022/2554 (DORA) or Directive (EU) 2022/2555 (NIS2), and, for all sectors, in the EU AI Act, which requires human oversight of high-risk systems to be assigned to people with competence, training and authority (Article 26). Sector supervisors also

publish their expectations on AI governance; one example, in insurance, is the EIOPA Opinion on AI governance and risk management (document 34).

> **Why it matters.** Where governance sits determines what happens on the day an important use case does not pass a gate. If the person who decides to stop reports to the person who has to deliver, nothing is ever stopped.

3.5 Six signals that the organisation is badly set up

1. Nobody can say within a week how many AI systems are in production.
2. The savings figures presented to the management committee have not been checked by anyone outside the team that produced them.
3. The second line only appears at the end, and can then only approve or veto.
4. No use case has ever been stopped on the merits; only for lack of budget or people.
5. The board sees product demonstrations and does not see portfolio, risks or validated value.
6. The same person proposes, approves and declares the benefit.

04

From the real organisation chart to the SEVEN-G bodies

SEVEN-G does not ask for the organisation chart to be changed: it asks for its bodies and roles to be **assigned** to people who already exist and for that to be written down (P38, with the rules of the bodies and the appointment of the AI lead).

SEVEN-G body or role	Who normally performs it	What is needed for it to count
Board of directors	The board itself, with an annual dedicated item and a quarterly item.	That it receives the pack in document 60, not a demonstration.
Board committee	The audit, risk or technology committee, whichever already exists.	A written mandate that includes AI (P38).
AI Committee	An extended management committee, or a dedicated committee with business, technology, risk, data and people.	Monthly frequency and real capacity to decide *gates* and retirements.

SEVEN-G body or role	Who normally performs it	What is needed for it to count
AI Office	Transformation, technology, the AI function or management control, depending on the company.	Dedicated time, method and access to the bodies; not building what it verifies.
AI Sponsor	Director of the unit that benefits.	Owns the benefit at the value realisation review (document 43).
AI Product Owner	Manager of the user unit or product manager.	Authority over scope and over adoption.
AI Technical Owner	Technology, data or the supplier, with an internal counterpart.	Never also the verifier.
AI Risk Owner and conformity	Risk, compliance, security and data protection.	Takes part from phase 0 and signs at G5 in Enterprise.
AI Auditor	Internal audit, with external support if needed.	Declared independence (P41).
Designated human overseer	A person in the user area who validates or interrupts the system.	Competence, training and **real authority to disagree** (document 50, section 4.2).

In small companies, document 30, section 11, and document 90, section 2.4, explain how to group bodies without losing the mandatory separations.

05

Where use cases come from

5.1 Nine sources

Demand for AI does not arise in one place. Recognising the source helps anticipate its bias and ask for what it lacks before it enters the portfolio.

#	Source	What it usually brings	Typical bias	What it almost always lacks
1	**Business units and operations**	Real problems, with their own process and data.	It brings the solution already chosen, not the problem.	Baseline and attribution method (document 40).
2	**Technology and data**	New capabilities that are now possible and cheap.	It looks for somewhere to apply what it already knows how to do.	A business sponsor who owns the benefit.

#	Source	What it usually brings	Typical bias	What it almost always lacks
3	**The product market**	AI features that appear inside software the company already uses.	They switch themselves on, with no decision and no inventory.	Inventory registration, regulatory classification and a conscious decision to enable them.
4	**Employees**	Individual use of general-purpose tools; the best small cases come from here.	It happens outside any control (unauthorised use).	A proposal channel, an acceptable use policy and AI literacy (document 31).
5	**Customers and channel**	Explicit requests and observed friction.	A one-off request is mistaken for aggregate demand.	Segment size and willingness to pay.
6	**Regulation, audit and incidents**	Obligations, findings and nonconformities that require change.	They are treated as cost and done to the minimum.	Fit with the portfolio; sometimes they enable value, not only control.
7	**Board and senior management**	Ambition challenges, sector comparisons, corporate transactions.	They arrive named after a technology and with no defined problem.	Translation into a falsifiable hypothesis and into sphere and ambition (document 10).
8	**Innovation ecosystem**	Proposals from suppliers, start-ups, venture capital, universities and innovation programmes.	A free proof of concept with no internal owner and no exit cost.	An internal sponsor, data that can be shared and an exit plan (document 36).
9	**The customer's agent**	Requests, purchases and claims that arrive already made by a customer's personal agent (35 §9.4).	It arrives as traffic, not as a proposal: nobody asked for it and nobody sponsors it.	A business owner who sets the posture by journey (13 §5.6) and answers for it.

Where use cases come from and how they enter

Nine sources, a single entry point to the register and the portfolio

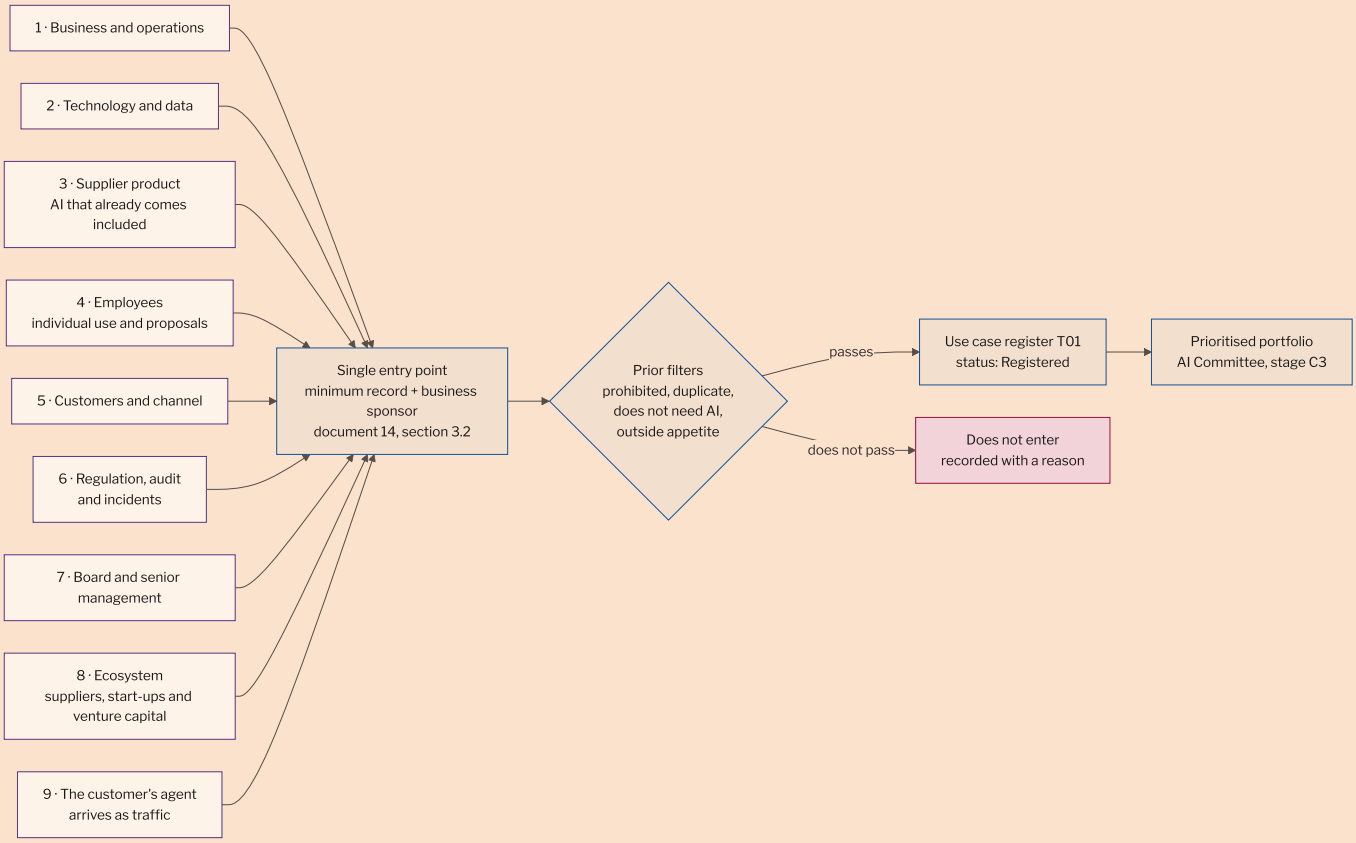

Source: SEVEN-G

5.2 There is a single entry point

Wherever it comes from, a use case enters through the same place: the **minimum entry record** in document 14, section 3.2 (plain-language description, sponsor and product owner, proposed sphere and ambition, order of magnitude of the value, investment up to G3, fit with the thesis and foreseeable Enterprise criteria), the **prior filters** in section 3.3 and the **entry windows** in section 3.4. Whatever does not enter is recorded with its reason: a "no" that is not recorded comes back three months later.

Sponsor rule. Every use case needs a business sponsor. A proposal from a supplier, from the ecosystem or from the technical area **with no internal sponsor does not enter**, however good it looks.

5.3 The three sources that cause most trouble

Source	Problem	How it is handled
AI that already comes included in a contracted product	It is enabled by default, with no decision, no inventory and sometimes with data that should not leave the perimeter.	Periodic review of supplier release notes; inventory registration and classification before enabling; clauses in document 36. It is a decision, not an update.

Source	Problem	How it is handled
Unauthorised use (shadow AI)	It happens anyway; banning it without an alternative simply hides it.	A usable corporate alternative, an acceptable use policy and an amnesty channel to surface uses: regularisation in document 14, section 11, and document 31.
The free proof of concept	With no apparent cost, it consumes data, time and credibility, and creates dependency.	Sponsor, success and stop criteria in writing before starting, bounded data and an exit cost estimated from day one.

> **Why it matters.** The three sources above share one feature: **they come in without passing through the entry point**. Most of a company's regulatory and security risk does not sit in the use cases in its portfolio, but in what never reached it.

06

How what has entered is governed

The route is the framework's own: the portfolio is managed as a funnel and each use case goes through phases with decision gates. What is added here is **who decides at each point** and with what degree of delegation.

Governance circuit of a use case

From proposal to operation, with who decides at each gate

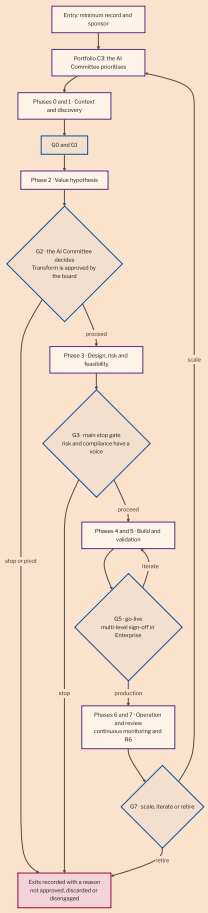

Source: SEVEN-G

Decision	Who decides (reference)	What cannot be delegated
Entry and priority in the portfolio	AI Committee, at its monthly review (document 14, section 7.1).	—
***Gates* for Lite use cases**	According to the delegation matrix in document 30, section 7.	Verification is done by someone other than whoever builds.
***Gates* for Enterprise use cases**	AI Committee, with the second line present.	Risk and compliance conformity.

Decision	Who decides (reference)	What cannot be delegated
Approval of Transform at G2 and of scaling at G7	Board (document 30, section 3.1).	Not delegated.
Acceptance of residual risk	The body matching the risk level (document 33).	Never whoever builds the use case.
Enterprise go-live	Multi-level sign-off at G5, with the second line.	The control sign-off.
Stop, pivot and retirement	AI Committee; the board for Transform use cases (document 14, section 10).	The reason is always recorded.
Exceptions to the policy	The body set out in P40, with an expiry date.	An exception with no end date is not an exception.

The role of the AI Office in this circuit is that of engine, not decision-maker: it prepares decisions, checks that the evidence exists before convening a gate, maintains the register and flags what has become stuck (stalled use cases, expired conditions, document 14, section 8).

07

How it is monitored and how the information reaches the board

7.1 The reporting chain

Each level receives less detail and more consequence. What does not change at any level is the **status of each figure**: validated, declared or estimated (document 40).

How AI information reaches the board

What each level sees, how often and with which tool

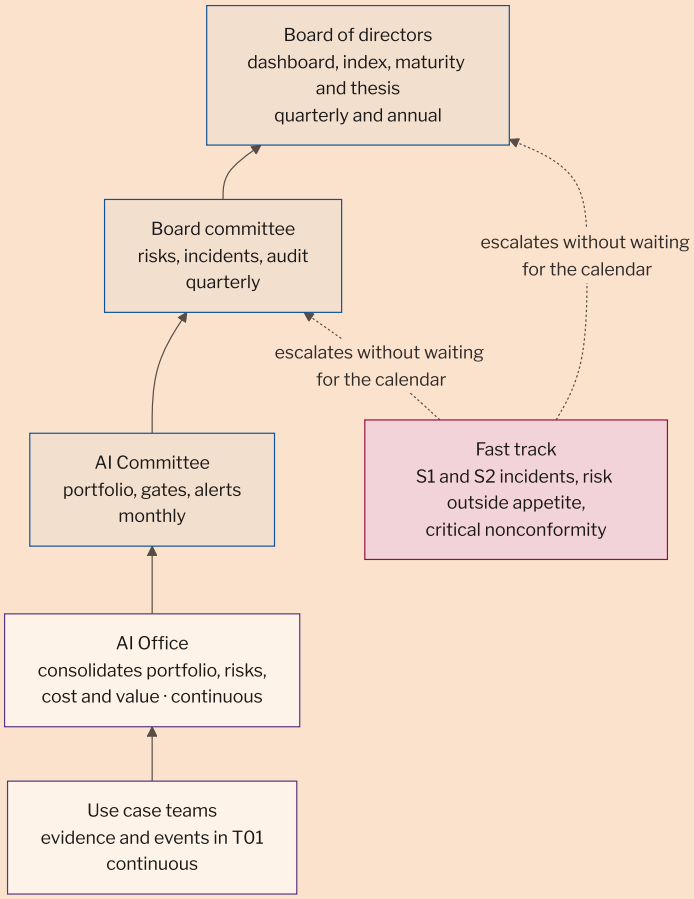

Source: SEVEN-G

7.2 What each level sees

Level	What it receives	How often	With what
Use case team	Its own evidence, the criteria for the next gate, open risks and its own figures.	Continuous	T01, T03, T06
AI Office	Full portfolio, time in phase, missing evidence, cost and consumption, alerts.	Continuous	T01, T11, T13
AI Committee	Prioritised portfolio, pending gate decisions, stalled use cases, High risks, retirements.	Monthly	T01, T06, P36
Board committee	Risks outside appetite, incidents, major and critical nonconformities, audit results, exceptions in force.	Quarterly	P42, document 37, document 38
Board	Portfolio funnel, validated versus declared value, maturity, transformation index, Transform bets and open recommendations.	Quarterly and annual	Dashboard T17, T14, T15, T18, pack P67 and document 60

7.3 Three rules for reporting to the board

1. **No figure without a status and a formula.** A declared saving and a validated one are not added in the same line (document 40).
2. **The board sees the whole funnel, not only what is going well**: the use cases that did not pass each stage and why, and those that have been disengaged after being in use. That is the part that teaches most.
3. **Alerts do not wait for the calendar.** Serious incidents, risks outside appetite and critical nonconformities escalate when they occur (document 37).

> **Why it matters.** A board that only sees demonstrations cannot exercise its oversight responsibility: it does not know what the portfolio costs, which risks the company has accepted or how much of the announced value has materialised.

08

What organisation comes next: the author's hypothesis

> **This is a reasoned hypothesis, not a forecast.** Nothing in this section is demonstrated, nor does it come from a market study. It is offered to support a management and board conversation, and each statement is accompanied by the signal that would allow one to check whether it is happening. The rules of the framework do not depend on this section being right.

8.1 Five forces pushing organisational change

Force	What changes
The cost of routine cognitive work falls	Advantage stops being about having more hands and becomes about defining the outcome well and controlling quality.
The unit of management shifts from the function to the process	If an end-to-end chain can run with few human interventions, splitting it across four functions stops making sense.

Force	What changes
Workers appear that are not people	Agents need identity, permissions, an owner, performance evaluation and offboarding: that is management, not only technology (document 35).
Knowledge becomes an operational asset	Whoever keeps what the AI reads up to date determines the quality of the answers: internal content moves from archive to infrastructure (document 51).
Regulation requires a human trail	The more a company automates, the more explicit human responsibility has to be: oversight, traceability and the ability to intervene (document 34).

8.2 What happens to each silo

Area	Hypothesis at three to five years	Early signal that it is happening
Business units	They stop asking for systems and start owning process outcomes; sponsorship becomes professional.	Directors' objectives linked to the benefit of the AI portfolio, not to the number of pilots.
Finance	From periodic close to continuous control; the controller designs automated controls and watches cost per unit of service.	AI consumption cost appears as a management line, allocated by case (document 42).
People	It manages a mixed workforce: staffing by capacity, not by posts; work redesign and oversight as a competence.	The workforce plan includes released capacity and its destination (document 50, section 6).
Shared services and back office	They shrink as executors and become operators of platforms and of exceptions.	The volume of routine tasks falls and the volume of exceptional cases handled by expert people rises.
Technology	From building applications to operating platforms, data, agents and behaviour quality control.	There is an internal catalogue of reusable components and agents with an owner.
Procurement	It contracts capabilities and outcomes, not licences; supplier dependency is managed as concentration risk.	The clauses in document 36 and the exit plan are a condition of purchase.
Risk and audit	They audit system behaviour, not only processes and documentary controls.	There is an annual AI audit plan with tests on real systems (document 38).
Customer service and marketing	The first point of contact is mostly automated, with people on exceptions and on the value relationship.	Usage drift and bias by segment are measured, not only response time (document 52).

8.3 Roles I expect to become established

Role	What it would do	Where it would come from
End-to-end process owner	Owns the complete outcome of a chain, with people and systems reporting to them.	Operations and business functions.
Agent engineer or manager	Designs, tests, deploys and retires agents; maintains their catalogue and their limits.	Technology and automation.
Non-human identity manager	Onboarding, permissions, credential rotation and offboarding of each agent and service.	Security and identity.
Consumption and cost-per-case controller	Consumption budget, reconciliation and cost per unit of service.	Management control.
Knowledge curator	Quality, currency and permissions of the sources the AI uses.	Documentation, quality and the areas themselves.
Professional human overseer	Validates, corrects and interrupts system decisions, with authority and with no incentive to accept them.	User areas, with specific training.
AI Auditor	Verifies evidence, behaviour and compliance with the framework.	Internal audit, with external support.

8.4 Four scenarios at five years

Scenario	What it consists of	When it is likely	Signal that it is heading there
1 · Reinforced silos	Each function uses AI within its own perimeter and the boundaries do not move.	Highly regulated sectors or ones with very separate processes.	Many use cases, none across areas.
2 · Central AI layer	A central platform and team serve all areas; the organisation chart does not change.	Companies with a strong corporate culture and a need for control.	The catalogue of shared components grows and areas stop contracting on their own.
3 · Process-based organisation with mixed teams	Work is organised into end-to-end chains, with people and agents, and outcome owners.	When the cost of coordinating between silos exceeds the cost of reorganising.	Process owners appear with budget and with outcome objectives.
4 · Governed network of agents	Much of the execution is carried out by coordinated agents; people set objectives, exceptions and control.	Only after high maturity in data, security and control.	There is an inventory of agents with an owner, limits and offboarding, just like an employee register.

The author's hypothesis is that most large European companies will travel through **scenario 2**, and that **scenario 3** will arrive process by process, not through general reorganisations: first one chain (customer onboarding, support, procurement, billing), then another.

8.5 What will not change

1. **Legal responsibility** remains with people and bodies, not with systems.
2. The **separation between building and controlling** becomes more important, not less.
3. It will still be necessary to **be able to stop**: the more automated the chain, the more valuable the switch.
4. Value is only demonstrated **in production and against a baseline**; no new organisation spares that proof.
5. Someone will have to **explain it to the board** in business language.

8.6 What a company that believes this hypothesis would do today

Four moves that are not regretted even if the future turns out differently: a **live inventory** of systems and agents; **consistent measurement** of value and of cost per case; **human oversight with authority** wherever decisions about people or money are made; and **a designated process owner** in the chain where coordination between areas hurts most.

09

The human organisation of an AI-first company

> **The author's proposal, not a framework rule.** This section puts in order what can be observed in consultancy publications, in announcements by companies from different sectors and in regulation, and proposes how to organise the people of a company that wants to be AI-first. It creates no rules or evidence: if it differs from documents 01, 30 or 50, those prevail. The sources consulted are in section 9.10.

9.1 What being AI-first means

There is no established definition. Some companies announced it in 2025 as a way of working (for example, the Duolingo post on "AI-first"). It is often used as a budget rule ("before asking for more

people, show that AI cannot do it"). That is a spending rule, not an organisation model: it does not say who owns what or where control sits.

Definition SEVEN-G proposes. A company is AI-first when **each process is designed on the assumption that AI does the work or proposes the decision, and people set the objective, decide on the exceptions and answer for the result**. AI stops being a project and becomes the normal way of working. Data, the platform and the controls are shared infrastructure.

The short formula is **AI by default, people where it matters**: AI moves the flow and a person is present at the moments that define the relationship with the customer, in decisions about people or money and in everything that cannot be undone.

What it is not:

- it is not "AI in everything";
- it is not replacing people as an objective;
- it is not having many use cases;
- it is not creating an AI function (section 9.5).

9.2 Six tests to tell whether a company is AI-first

The six tests serve as indicators. Each one rests on something SEVEN-G already requires or measures.

#	Test	How it is checked in SEVEN-G
1	**A single end-to-end owner.** Each key process has one person who answers for its complete result, with a budget, even when the process crosses several functions.	Business sponsor and business benefit owner (document 01, section 8; document 43).
2	**The default question is "why not with AI?"**, with control proportionate to risk, not the same funnel for everything.	Express, Lite or Enterprise intensity (document 01, section 9) and the proportionality principle (document 01, section 3).
3	**Every agent has a named human owner**, appears in the inventory and counts in the workforce plan.	AI system inventory (T02) and non-human identities (P54); workforce plan in document 50.
4	**Controls are built into the platform** and apply before production, not afterwards.	Feasibility gate (G3) and go-live (G5); security controls in document 35.
5	**Cost is measured per outcome** (people, agents and consumption), not per head.	Full cost per case (document 42) and validated value (document 40).
6	**Decisions leave a trail** (what AI proposed, what was decided and what happened) and the company learns from it.	Logging and monitoring (document 52); board decision register (document 62).

> **Why it matters.** "AI-first" is often used as a slogan. With six verifiable tests, the board can ask how many are met today and when the rest will be. Test 4 is usually the most urgent: a case in production without closed control is an open risk.

9.3 Why a function-based organisation chart holds things back

The classic organisation chart divides the company into vertical functions: business, technology, operations, risk and people. Useful AI lives in processes that cross those functions. For example, onboarding a business customer starts in sales and goes through risk, contracting, billing and customer service, and it is touched by business, operations, technology, security and compliance.

If nobody owns the whole process, each function optimises its own stretch, the case stays a pilot and control arrives late.

A use case today

The same case depends on five areas and none owns the whole process

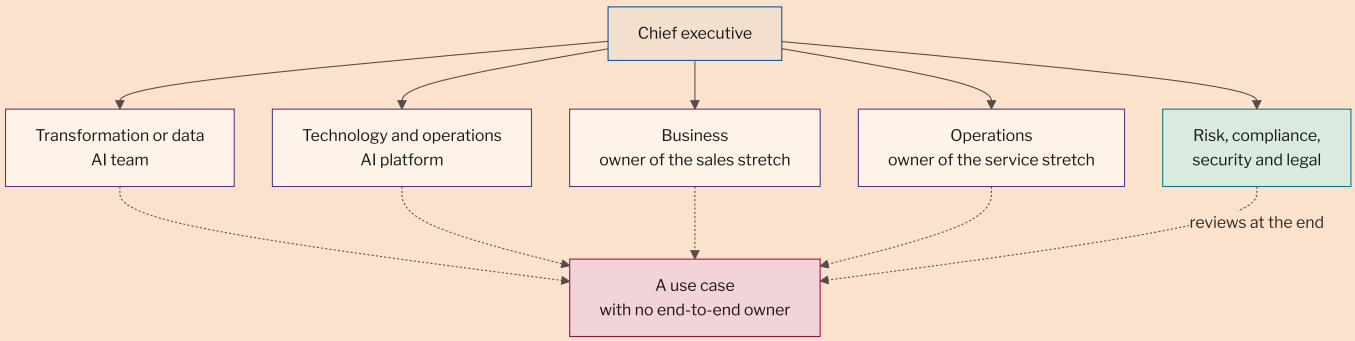

Source: SEVEN-G

Where a case gets stuck today and where it fails. Eight recurring points:

#	Type	Where	What happens	What solves it in the proposed model
1	Bottleneck	Entry	The case waits its turn in a committee that prioritises single cases, not processes.	The chain owner prioritises within their budget (section 9.6, option C).
2	Bottleneck	Control	Risk questionnaires pile up and are reviewed one by one.	Lanes by risk level, with controls built into the platform.
3	Bottleneck	Integration	The case needs systems from another function that has other priorities.	A single AI and data factory (option A).

#	Type	Where	What happens	What solves it in the proposed model
4	Bottleneck	Full cycle	Months pass between the idea and production with no decision date.	Gates with deadlines and funding in tranches (document 14, section 6).
5	Problem	Design	The idea solves one stretch, not the process: local improvement, same overall result.	Redesign of the whole process by its owner.
6	Problem	Build	Each area builds its part with its own tool and supplier.	Shared platform and component catalogue.
7	Problem	Production	The system goes live with control only half closed.	Go-live (G5) with a control sign-off.
8	Problem	Value	Nobody measures the outcome: activity is reported, not benefit.	Baseline and benefits realisation (documents 40 and 43).

9.4 What can be observed elsewhere

Who	What it does or proposes	What it means for your company
Accenture, consultancy ("When everyone has AI, your operating model will set you apart ", 2026)	Organise the company into five to eight end-to-end value streams. Each stream has accountability, decision rights and budget. Every agent has a named human owner.	Give each key process an owner with their own budget. Departments still exist, but the person who answers for the result of the process is that owner.
BCG, consultancy ("AI-First Enterprise Operations: Reinventing the Operating System of Work ", 2026)	A process owner who answers for the end-to-end result, and a central "factory" that concentrates funding, guardrails and priorities.	Appoint an owner for the whole process and create a single central team that provides the platform and the rules. Adding an assistant to each job improves little; the large improvement comes from redesigning the whole process.
Microsoft, technology (Microsoft Work Trend Index 2025 report ; a cross-sector study)	Teams made up of people and agents. A new role for everyone: someone who builds agents, delegates work to them and manages them. Work is organised around outcomes and a new measure appears: how many agents there are per person.	Prepare people to direct agents, not only to use them. Decide in each process how many agents one person can supervise well.
BBVA, banking (BBVA note on its AI Transformation area , 2026)	Creates, at the top level of the organisation, a global area that joins the data area and the technical capabilities to create, deploy and manage agents across the company.	AI, data and the platform sit together in a single top-level area. It is the "single factory" of section 9.6: it gives capacity to everyone, but it does not replace the owner of each process.

Who	What it does or proposes	What it means for your company
Moderna, pharmaceuticals (Moderna leadership team)	A single executive brings together the people function and digital technology.	If agents do part of the work, the workforce plan must count people and agents together. For that, people and technology have to plan jointly.
Allianz, insurance (Allianz article on Project Nemo, 2025)	Seven agents, each with one task, handle simple claims. The payout is always decided by an accountable person.	The whole process is redesigned with agents, but the decision that affects the customer and the money still belongs to a named person.
McKinsey, consultancy (three 2026 articles on adoption horizons, organisation and operating model)	There is no single operating model. The most advanced organisations organise teams by end-to-end outcome, move budget frequently and give one executive real authority over funding and talent. They propose drawing, for each outcome, what people do and what agents do.	Do not copy an organisation chart. Decide structure, governance, processes, talent and rewards consistently, and review those decisions every quarter. Section 9.11 develops this.
MIT NANDA, university research ("The GenAI Divide" report, 2025; preliminary findings)	General-purpose tools are widely adopted, but custom systems rarely move from pilot to production. The cases with the best results start with line managers, not in a central lab.	Give whoever runs the process the authority to choose and deploy, and keep them accountable for the outcome. Read its figures with caution (section 9.11).
Regulation, all sectors (EU AI Act, Articles 4 and 26)	Every company that uses AI must ensure the AI literacy of its staff. Whoever uses a high-risk system must assign human oversight to people with competence, training and authority.	The law does not ask for a new role. It asks that each system has named people with the training and authority to oversee it.

What recurs:

1. the end-to-end process is the unit of design;
2. there is a small central piece that provides platform, standards and controls;
3. business and technology share objectives;
4. agents are managed as part of the workforce, with a human owner;
5. none of the references creates an AI function separate from technology and data.
6. there is no single model: what sets the most advanced organisations apart is the consistency of their decisions and how well they carry them out.

Caution. Almost all of the material comes from consultancies with a commercial interest. Their savings figures come from their own clients and SEVEN-G does not repeat them. Company announcements say how they are organised, not whether it works for them.

9.5 Is a chief AI officer needed? In general, no

Section 3.3 gives four questions to decide whether a central function is needed. That function is needed; **an AI function as a vertical silo, in general, is not**:

- **it is a vertical silo for something horizontal:** the chief AI officer ends up owning the use cases and the processes still have no owner;
- **it creates one more pole** between the areas that already share AI, instead of resolving the split;
- **the regulator does not ask for it:** it asks for clear accountabilities (section 9.4);
- **if it also builds and controls,** it breaks the separation of functions (section 2).

Who answers to the board. The chief executive, with three responsibilities divided as follows:

Responsibility	Who exercises it	What they answer for
Value	The value chain owners	The result of their process, with budget and agents.
Capacity	The head of the AI and data factory	Platform, data, model operations and inventory.
Control	The second line	Risk, compliance, security and legal, with controls built into the platform.

If the company wants a title, it should be a coordination role within the area that holds the factory, not one more function. The AI Office in document 30 keeps doing its work on method, portfolio and reporting.

9.6 Five organisation options

Option	What it consists of	For	Against	Fit
A · A single AI and data factory	The AI teams, the platform, model operations and the inventory are brought into one unit, with a single funnel. Areas still request the cases.	Small change; ends AI being spread across several homes.	Nobody owns the whole process; AI is still "something technology does".	First step

Option	What it consists of	For	Against	Fit
B · Hub and spoke	The factory from A as the hub and an AI lead in each function and subsidiary, with a functional line to the hub.	Respects a federated group.	A hard matrix to staff; local leads end up absorbed by their area; there is still no process owner.	For subsidiaries and distant units
C · Value chains	Five to seven end-to-end chains, each with a business lead and a technology partner who share objectives, a budget in tranches and their own agents. Functions become centres of competence (standards, experts, careers).	The only option that resolves horizontality.	Moves power between functions; needs the chief executive behind it; done all at once, it can paralyse operations.	Destination
D · Temporary transformation office	A small office reporting to the chief executive, with a closing date (24 to 36 months), that manages the portfolio, funds in tranches and redesigns processes with the areas.	Speed and clear sponsorship without a permanent silo.	Risk that it becomes permanent or acts as a parallel organisation.	Engine of change
E · Mixed workforce	People and technology plan the workforce of people and agents together; new roles are defined (agent owner, agent supervisor, chain owner) and variable pay is linked to AI and security objectives.	Addresses the human side: talent, roles and careers.	Not enough on its own.	Always

Option	Resolves horizontality	Change and risk	Fit
A · Single factory	No	Low	First step
B · Hub and spoke	A little	Medium	For subsidiaries
C · Value chains	Yes	High	Destination
D · Temporary office	While it lasts	Medium	Engine
E · Mixed workforce	No	Low	Always

Relation to the patterns in section 3.2: A and B are variants of pattern C (central office or function); option C in this section is scenario 3 of section 8.4 made real chain by chain.

> **Why it matters.** Choosing only the factory (A) is easy to agree and to check, but it leaves processes without an owner. Choosing value chains (C) all at once solves the most and carries the most risk. The phased combination (section 9.7) takes the best of each option.

9.7 Recommendation: a phased combination

Destination C, built on A, started by D, with B for subsidiaries and E from day one. In one sentence: processes have an owner, AI has a single home and control sits inside the platform.

The target model

Value chains with an owning pair on a single factory, with control inside the platform

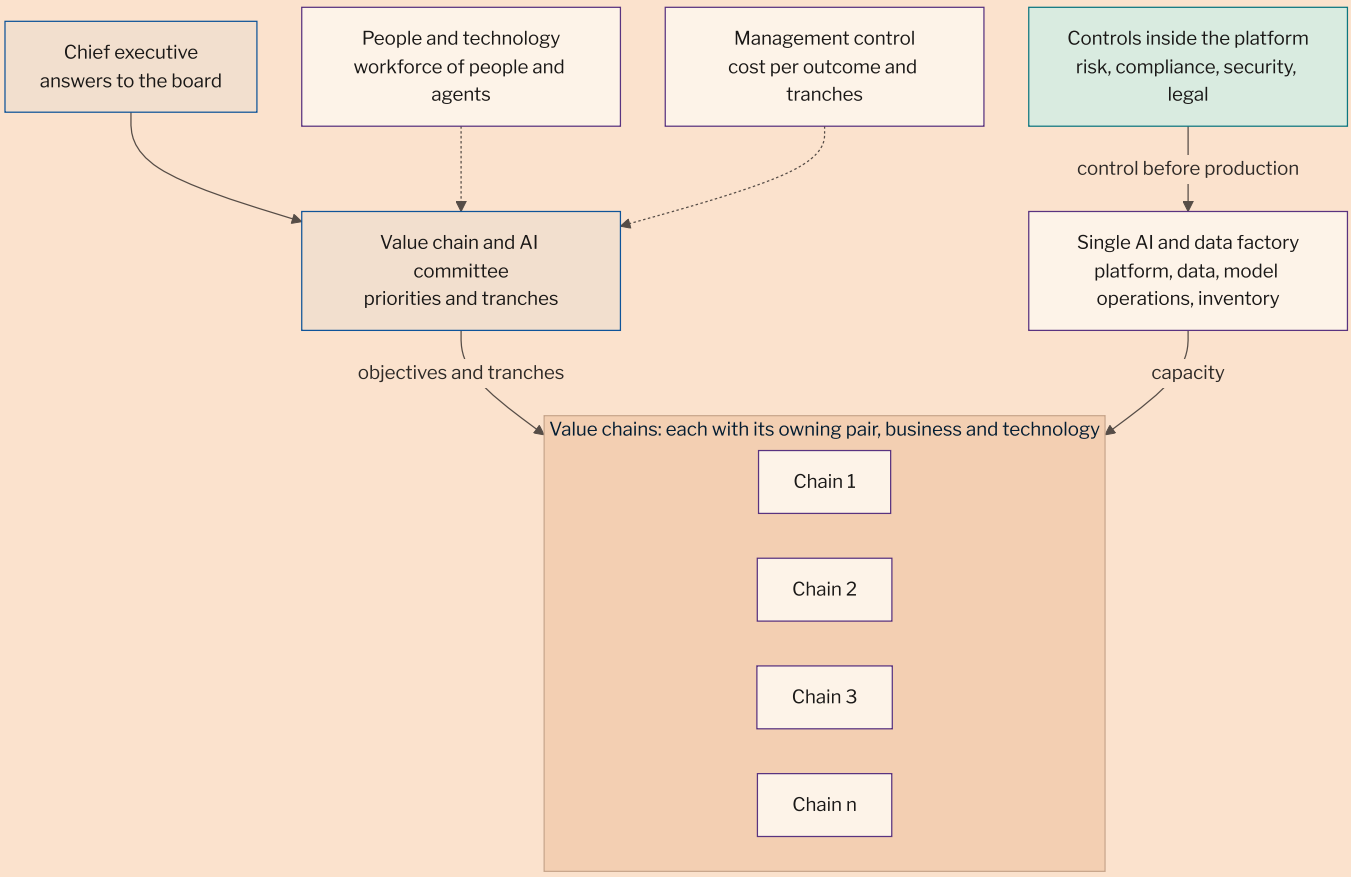

Source: SEVEN-G

Each chain has an **owning pair**: a business lead and a technology partner with the same objectives. The chains hold the result and the budget; the factory, the capacity; the second line, the control. The committee decides priorities and tranches.

How a case travels in the target model

From the redesign proposed by the chain owner to the review against the value curve

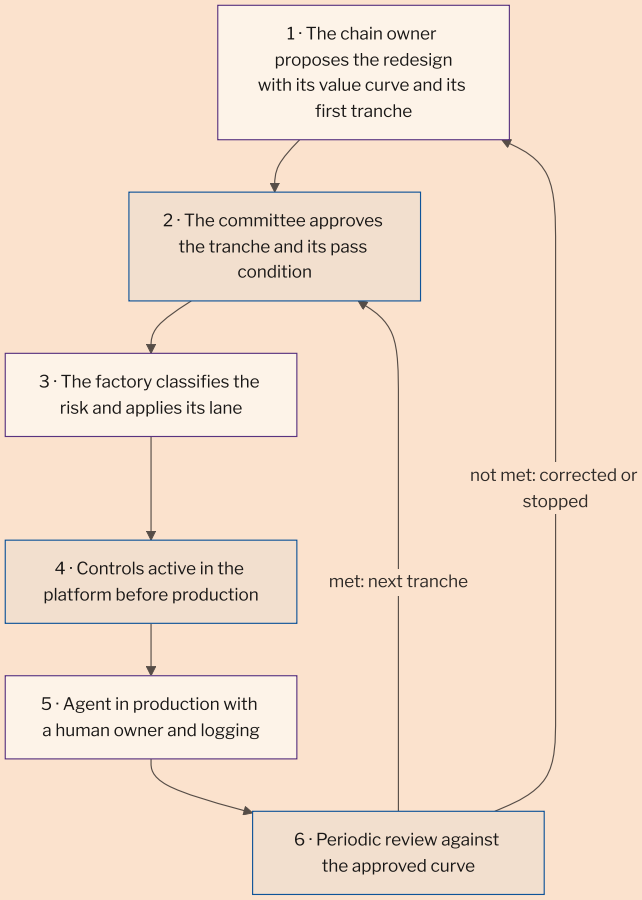

Source: SEVEN-G

Control comes before production and money is released in tranches according to real value (document 14, section 6.2; document 43). The sequence of phases and gates in document 01 does not change: what changes is who owns the case.

Step by step:

Stage	What is done	Why	If it fails
0 · Decide and appoint (first two months)	The chief executive appoints the temporary office and two pilot pairs in the two chains where AI already works or where coordination hurts most. A single inventory tool is chosen.	Without the chief executive's sponsorship, destination C cannot hold.	The company falls back to a function-based model and the problem is postponed.

Stage	What is done	Why	If it fails
1 · A single factory (first half-year)	The AI teams and the platform are brought under one head, with one funnel, one inventory and lanes by risk; whatever is in production without closed control is regularised (document 14, section 11).	The chains need a single home to ask.	The pilots ask two places and the split repeats itself.
1 · Pilots with an owner (first half-year, in parallel)	The two pairs receive objectives, their chain's AI budget in tranches and ownership of their agents. People opens the mixed workforce plan and the new roles. Management control measures cost per outcome.	Prove it with data before extending it.	It shows in the periodic review: the curve is not met and the next tranche is corrected or stopped.
2 · Extend (second half-year and the following year)	The other chains are added; the AI committee becomes the value chain and AI committee, chaired by the chief executive; variable pay is linked to the chain's objectives; subsidiaries join with model B.	Extend what has worked.	Matrix conflict between chain and function; mitigated with shared objectives and a clear arbiter (the committee).
3 · Close the temporary office (24 to 36 months)	Its functions move to the chains and the factory.	Avoid a permanent silo.	If the office does not close, an AI function has been created through the back door.

Risks that remain in the proposed model and how they are covered:

#	Risk	How it is covered
1	The chain promises value without a baseline.	Baseline and attribution method before the first tranche (document 40).
2	The committee becomes another bottleneck.	Delegation by intensity (document 30, section 7); the committee decides tranches, not single cases.
3	The factory is saturated and the chains wait.	Factory capacity in the enablement budget (document 13); committee priorities.
4	Control arrives at the end again.	Controls inside the platform and the feasibility gate (G3) with the second line present.
5	New roles with no definition or career path.	Mixed workforce plan and roles in document 50; written appointments (P38).
6	A tranche's pass condition is soft.	Measurable pass condition and review against the approved curve (document 43).

9.8 Implications of deciding only the factory or the phased combination

The two options are not mutually exclusive: the single factory is the first step of the combination. What is really decided is **what is written down as the destination**.

	Phased combination	Single factory only (A)
What it solves	AI spread across several homes and no process having an owner.	Only the spread: one funnel, one inventory, one head.
Type of value	Redesigns the whole process, where the large savings and a different service are.	Case-by-case improvements, better ordered.
Power	Business gains its chain's AI budget; the AI teams become suppliers to the chains.	Only which area keeps the factory is decided.
Sponsorship	The chief executive, visible and sustained.	Two functions agreeing is enough.
Timeframe	First data in about six months; extended model in two years.	Working in three to six months.
Operational risk	Higher if it coincides with other large changes; reduced by starting with two chains where AI already works.	Low.
Control	Every agent has a business owner and controls come before production.	The funnel is put in order, but control may still arrive at the end.
People	New roles, pay linked to objectives and more attractive to talent.	Almost nothing changes; AI remains a matter for specialists.
Unintended effect	Matrix conflicts between chain and function; an arbiter is needed.	The head of the factory ends up, in practice, as the chief AI officer nobody wanted.

Middle way: agree the single factory now and test two value chains with data before closing the target model. It costs little and keeps the door open.

9.9 What exists and what is missing: checklist

Piece	Question	Where it rests in SEVEN-G
AI policy and bodies	Do the bodies decide on processes and budget, or only on cases?	Document 30; P38
AI platform	Is there a single home for AI, with one funnel and one inventory?	Document 14; AI system inventory (T02)

Piece	Question	Where it rests in SEVEN-G
Lanes by risk	Do controls work before production?	Document 01, section 9; documents 33 and 35
End-to-end process owner	Are there owning pairs with a budget?	Document 01, section 8; document 43
Human owner of each agent	Does each agent have an owner, onboarding and offboarding in the inventory?	P54; document 35
Workforce plan for people and agents	Does People also plan the agents?	Document 50, section 6
Cost per outcome and tranches	Does management control measure cost per outcome and release funds in tranches?	Documents 42 and 43; document 14, section 6.2
Decision trail	Is it recorded what AI proposed, what was decided and what happened?	Document 52; document 62

9.10 Sources consulted and cautions

External sources checked on their original site on 10-11-2026 and 13-11-2026 and entered in the reference register (D41, D114):

Source	Type	What is taken
Accenture, "When everyone has AI, your operating model will set you apart" (06-10-2026)	Consultancy publication	End-to-end value streams with budget; a human owner for each agent.
BCG, "AI-First Enterprise Operations: Reinventing the Operating System of Work" (2026)	Consultancy publication	End-to-end process owner and a central factory for funding, guardrails and priorities.
Microsoft Work Trend Index 2025 report (23-04-2025)	Study by a technology company, across sectors	Teams of people and agents; the role of whoever directs agents; work organised around outcomes.
BBVA note on its AI Transformation area (28-05-2026)	Company announcement (banking)	Top-level area that joins data and the technical capability for agents.
Moderna leadership team	Company page (pharmaceuticals)	People and digital technology under one executive.
Allianz article on Project Nemo (03-11-2025)	Company announcement (insurance)	Task-specific agents and payout decisions always made by a person.

Source	Type	What is taken
Duolingo post on "AI-first" (2025)	Company announcement (digital education)	Use of the term "AI-first" as a way of working.
EU AI Act, Articles 4 and 26	European Union law	AI literacy and human oversight with competence, training and authority.
McKinsey, "The agentic organization: contours of the next paradigm for the AI era" (mckinsey.com, 26-09-2025)	Consultancy publication	Small teams of people, owners of an end-to-end outcome, that supervise the agents of a whole process. The person keeps final accountability. Human oversight capacity limits how far it can scale. Its figures are not repeated.
McKinsey, "From adoption to impact: Three horizons of AI transformation" (mckinsey.com, 8 July 2026)	Consultancy article with a survey of 750 employees and executives	Three adoption horizons: enablement, automation and reinvention. Organisational readiness weighs almost twice as much as that of people. 11% of executives place their organisation in reinvention. The answers are each respondent's opinion about their organisation.
McKinsey, "The key to AI value is hiding in plain sight: Your operating model" (mckinsey.com, 9 September 2026)	Consultancy article with a survey of over 700 executives	There is no best operating model. Twelve design decisions and the ones most frequently chosen by the most advanced organisations. Automating without redesigning gives a worse result in clarity and in capabilities. The improvement figure of one of its clients is not repeated.
McKinsey, "AI is changing work. Now it has to change the organization" (mckinsey.com, 5 October 2026)	Consultancy article	Orchestration map of people and agents by outcome. Explicit decision on the destination of released capacity. How newcomers learn when AI does the entry-level work. An executive with authority and a stable team that redesigns, builds and controls at the same time.
MIT NANDA, "The GenAI Divide: State of AI in Business 2025" (media.mit.edu, July 2025)	Report of a university project; preliminary findings from 52 interviews, 153 executive surveys and over 300 public cases	Custom systems rarely move from pilot to production. The main cause is that they do not learn from use or adapt to the process. The cases with the best results start with line managers. It was read in full in the PDF that the author obtained from the publisher. Its figure for organisations with no return is not repeated: the report applies it to organisations and to pilots at the same time.
"The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise" (nber.org, NBER Working Paper 33641, April 2025; also Working Paper 25-043 of Harvard Business School)	Academic study: field experiment with 776 professionals at a consumer goods company	One person with AI performed like a pair without AI. Pairs with AI produced more solutions in the top 10%. With AI, the proposals of technical and commercial profiles became more alike. It was read in full. It is a working paper, not peer reviewed: one company, one ideation task and a single day. The effect on silos is shown in a chart, without a statistical test.

The McKinsey publications were read on their official site with an ordinary browser (the first on 13-11-2026 and the 2026 ones on 19-11-2026), because the site does not respond to automated reading. A McKinsey publication from 2021 on the platform model in banking was also consulted; it is not cited because it is old and predates agents.

The MIT NANDA report and the study "The Cybernetic Teammate" were read in full on 19-11-2026. Before that, the first was listed as not checked because its publisher delivers it after a form, and only the abstract of the second had been read. No source in this section remains unchecked.

The examples are chosen from different sectors so that none weighs more than another; for that reason the announcements of two other insurance groups that were consulted are not cited here, nor a 2026 report on AI in the insurance sector, which is cited in the sector gallery.

Cautions. Consultancies' savings figures are not repeated because they come from their own clients. Company announcements describe how they are organised; they do not prove that it works for them. The options and the recommendation are the author's proposal for a management or board conversation.

9.11 What the 2026 sources add: operating model, talent and transformation engine

Three McKinsey articles from 2026, the MIT NANDA report of 2025 and the experiment "The Cybernetic Teammate" extend what section 9.4 gathers. This section says what they add and how to read it with SEVEN-G. It remains a proposal for a management conversation, not a rule.

1. There is no single operating model. In McKinsey's survey of over 700 executives, the organisations that have most redesigned their work with AI do not share a structure. They share a consistent combination of decisions and they carry it out well. These are the ones chosen most frequently:

Decision	What they choose most frequently	How it reads in SEVEN-G
Structure	Cross-functional teams organised by product, by customer journey or by end-to-end process, accountable for the outcome.	End-to-end value chain (section 9.6, option C).
Governance	Quarterly objectives or a forecast that is continuously redone, instead of three-to-five-year plans. What works is scaled up and what does not is stopped.	Decision gates with stopping as a valid outcome and funding in tranches (document 14, section 6.2).
Processes	Standardised processes, with metrics and with the human decision point defined.	Human oversight (P17) and monitoring (document 52).
Leadership	Fast decisions, taken close to where the work is done.	Who decides at each gate (document 30).
Talent	Management by capabilities, not by fixed jobs.	Capability plan (document 50, section 5).

Decision	What they choose most frequently	How it reads in SEVEN-G
Rewards	Mixed: team outcome and each person's contribution.	SEVEN-G does not regulate it. It is a decision for the people function.
External partners	An explicit decision on what is built and what is bought.	Third parties and suppliers (document 36).

A warning from the same survey: organisations that automate the workflows they already have, without redesigning them, report **less** clarity and fewer capabilities than those that only gave tools to their employees. Automating the current organisation chart is not a safe intermediate step (section 9.3).

2. From the organisation chart to the orchestration map. McKinsey proposes drawing, for each outcome, which people and which agents take part, what each does, which handovers disappear and where a person decides. The agent's autonomy, its access, its oversight and whom it escalates to are part of the design of the work; they are not added at the end. In SEVEN-G those pieces already exist: agent inventory with a human owner (P54), autonomy levels A0 to A3 (document 35) and oversight design (P17). The orchestration map brings them together in one drawing per process.

- In one example in the article, product teams go from eight or ten people to four to six with agents. It is one case, not a general measure.
- Fewer levels and fewer coordinators only give value if the coordination that disappears is replaced by better information, clear decision rights and more judgement from people.

3. A transformation engine with authority. Transformation crosses business, technology, people and risk, and no current function has the full mandate. McKinsey asks for an executive with real ability to move funding and scarce talent and to resolve priorities; it says that the title matters less than those levers. It matches section 9.5. It also proposes a stable team that takes a business outcome, redesigns the flow of people and agents, builds the agents and their controls, puts it into operation and leaves reusable pieces. Business, technology, risk, compliance, security, legal and people work together from day one. This is the AI and data factory of section 9.6 working with the owner of the value chain. The

SEVEN-G gates already require that evaluation, permissions, monitoring and the accountable person are defined before production.

4. What the research says, with its limits.

Source	What it observes	Limits
MIT NANDA, "The GenAI Divide" (2025)	General-purpose tools are tried and widely adopted. Custom systems rarely reach production: in its interviews, out of every 100 organisations, 60 evaluated them, 20 piloted them and 5 put them into production. The main cause is neither technology nor regulation: it is tools that do not learn from use or adapt to the process. The cases with the best results start with line managers who answer to management. Pilots done with a supplier reached deployment twice as often as those built in-house. Employees use personal AI accounts in the great majority of companies.	Preliminary findings. Figures declared by interviewees, with inconsistencies between the text and the charts. The report itself warns that it does not prove that buying causes the better result. The publisher promotes an agent infrastructure. It presents the use of personal accounts as an opportunity and does not analyse its risk: in SEVEN-G it is unauthorised use (section 5.3).
"The Cybernetic Teammate" (2025)	In a real product challenge, one person with AI performed like a pair without AI. Pairs with AI produced more solutions in the top 10%. With AI, the proposals of technical and commercial people became more alike.	One company, one ideation task and one day. It does not show that AI replaces teams or that the effect lasts. Four of the authors work at the company studied.

What changes in the recommendation of section 9.7. Nothing in substance: the sources reinforce the end-to-end process, the small central piece and business accountability. They add three clarifications:

1. **Review frequently.** The organisation is adjusted every quarter with portfolio data; it is not designed once.
2. **Draw the split.** Each value chain has its map of people and agents, with the human decision point marked.
3. **Decide what to buy.** Build in-house only what gives an advantage. Buy the rest, with the requirements of document 36.

> **Why it matters.** Most companies give tools to their employees and expect value to appear. The 2026 sources agree that it does not: value arrives when workflows, roles and who decides change. That is an organisation decision, and only management can take it.

10

Questions for a useful conversation

For the board, the management committee or a working session. They have no right answer; they either have an answer or they do not, and that in itself is information.

On who does what

1. Who performs each of the six functions in section 2 today, and which are empty?
2. Is anyone performing both build and control? Since when?
3. If an important use case has to be stopped tomorrow, who signs that decision?

On demand

4. Which of the nine sources do our current use cases come from? Is there any source that never appears?
5. How many AI systems have been enabled in products we already had under contract, with no express decision?
6. What do we do today about the AI use that employees already practise on their own?

On governance

7. How many use cases have been stopped on the merits in the last twelve months? If none, why?
8. Does the second line take part from phase 0 or does it appear at the final gate?
9. Which exceptions to the policy are in force and when do they expire?

On reporting

10. How much of the value we present is validated, how much declared and how much estimated?
11. Does the board see the use cases that did not make it, and why?
12. How much does the AI portfolio cost per year, including consumption and recurring cost?

On the future

13. Which of the four scenarios in section 8.4 are we in, and which do we want to be in?
14. Which end-to-end chain of work suffers most from the separation between areas?
15. If in three years part of the work is done by agents, who onboards them, who owns them and who offboards them?

On being AI-first

16. How many of the six tests in section 9.2 do we meet today?
17. Which process has an end-to-end owner today, with a budget?
18. If the company's AI is spread across several areas, which of the five options in section 9.6 brings us closest to the destination without stopping operations?

11

How it is set up in practice

The implementation sequence is in document 90. In organisational terms, the order that avoids most blockages is:

1. **Census and inventory** of what already exists, including what came in without passing through the entry point (document 32; document 14, section 11).
2. **Assign the six functions** by name, even if they are concentrated in few people (P38).
3. **Open the single entry point** and publish the minimum record, so that demand stops coming in through the sides (document 14, section 3).
4. **Settle the reporting to the board** before there are many use cases: it is easier to start measuring with five than with fifty (document 60).
5. **Review the structure after six months** with your own data: how many use cases, from which sources, how many stopped, how much cost. Structure is better decided with that information than with a comparison of other companies' organisation charts.

12

Associated tools and templates

Name	Use in this document
Use case register (T01)	Single entry point, status of each use case and origin of all later information.
AI system inventory (T02)	Census of what already exists, including AI embedded in contracted products.
Risk matrix and register (T06)	Risks by use case and their acceptance by the appropriate body.
Transformation index and maturity diagnosis (T14 · T15)	Overall reading for the board and for the annual review.
Board dashboard and recommendations register (T17 · T18)	Periodic reporting to the board and follow-up of its agreements.
Rules of the bodies and appointment of the AI lead (P38)	Written assignment of bodies and functions.
Agenda and minutes (P39)	Operation of the AI Committee and of the board committee.

Name	Use in this document
Exceptions (P40)	Exceptions with body, reason and expiry.
Independence and conflict declarations (P41)	Separation of duties and audit.
Quarterly second line report (P42)	Risk and compliance reporting to the board committee.
Portfolio plan (P36)	Prioritisation, traffic light and regularisation.
Quarterly board pack (P67)	Content of the quarterly session.

13

Related documents

Document	Relationship
01 · Foundational methodology	Use case roles, separation of duties, phases, gates and intensity.
13 · AI thesis, ambition and risk appetite	What the board approves, which gives content to the direction function.
14 · Portfolio management	Entry channels, minimum record, filters, prioritisation and retirement.
20 and 21 · Phases and *gate* criteria	What is done in each phase and what is verified at each gate.
30 · Governance model	Bodies, mandates, delegation matrix, escalation and three lines. Reference document for this one.
32, 33 and 34 · Inventory, risk and regulation	Classification, risks and obligations that determine control.
38 · AI audit framework	Independent verification function.
40 and 42 · Value measurement and costs	Status of the figures and cost per case reported to the board.
50 · People and organisation	Effect on work, new roles, human oversight and released capacity.
60, 61 and 62 · Board	Pack, conversation and register of decisions and recommendations.
90 · Implementation guide	Start-up sequence and minimum route in Lite scope.
94 · Obligation matrix	Level of each document and layered reading.

14

Version control

Version	Date	Changes
0.1	22-09-2026	First version. Six functions someone must perform and their separations; organisation patterns observed in the market and where to place AI governance; mapping between the real organisation chart and the bodies in document 30; eight sources of use cases with their bias and their single entry point; governance circuit with who decides at each gate; reporting chain to the board; forward-looking section marked as a hypothesis, with four scenarios and their signals; fifteen questions for a management or board conversation. Explanatory document: it adds no rules or evidence.
0.8	19-11-2026	New section 9.11, "What the 2026 sources add: operating model, talent and transformation engine". Section 9.4: McKinsey and MIT NANDA rows. Section 9.10: three McKinsey articles from 2026; the MIT NANDA report and the study "The Cybernetic Teammate" become read in full, with their cautions (D194).
0.7	16-11-2026	Section 5.1: ninth source of use cases, "The customer's agent", and "eight" becomes "nine" in the title, the chart and question 4 of section 10 (D185).
0.6	13-11-2026	Section 9.10: the Harvard Business School study "The Cybernetic Teammate" is added (field experiment; AI breaks down functional silos) and the 2021 McKinsey publication on banking is removed (old and predates agents).
0.5	13-11-2026	Section 9.10: the two McKinsey publications become checked sources (read on their official site with an ordinary browser) and the text says what is taken from each. The MIT NANDA report is still not checked: its publisher delivers it after a form.
0.4	13-11-2026	Section 9.10: the internal Shopify message (a leak of an internal document) and the 2018 DBS annual report (it is old) are removed from the consulted sources.
0.3	13-11-2026	Section 9.4: examples from different sectors (consultancy, technology, banking, pharmaceuticals, insurance and general regulation) and a "What it means for your company" column with full explanations. Sections 3.4, 8.4 and 9.10 adjusted so that no sector weighs more.
0.2	10-11-2026	New section 9, "The human organisation of an AI-first company": proposed definition and six tests; why a function-based organisation chart holds things back; eight points where a case gets stuck or fails; what can be observed elsewhere, with checked sources; why, in general, an AI function is not needed; five options; recommendation of a phased combination with the target model, how a case travels, the step-by-step plan and the remaining risks; implications of each decision; checklist; sources and cautions. Questions 16 to 18. Sections 9 to 13 become 10 to 14. The author's proposal: it adds no rules or evidence.

External sources cited

Links to the official or primary source, checked before publication. Always verify the version in force on the date of use.

Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS 2 Directive) · https://eur-lex.europa.eu/eli/dir/2022/2555/oj/eng · Parlamento Europeo y Consejo de la UE · 14-12-2022

Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA) · https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng · Parlamento Europeo y Consejo de la UE · 14-12-2022

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) · https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng · Parlamento Europeo y Consejo de la UE · 13-06-2024

EIOPA Opinion on Artificial Intelligence governance and risk management (EIOPA-BoS-25-360), 6 August 2025 · https://www.eiopa.europa.eu/publications/opinion-artificial-intelligence-governance-and-risk-management_en · Autoridad Europea de Seguros y Pensiones de Jubilación (EIOPA) · 06-08-2025

Allianz: How Agentic AI Is Transforming Insurance Claims Processing (Project Nemo) · https://www.allianz.com/en/mediacenter/news/articles/251103-when-the-storm-clears-so-should-the-claim-queue.html · Allianz SE · 03-11-2025

BBVA Accelerates Its Artificial Intelligence Strategy With A Global Area: AI Transformation · https://www.bbva.com/en/innovation/bbva-accelerates-its-artificial-intelligence-strategy-with-a-global-area-ai-transformation/ · BBVA · 28-05-2026

Microsoft: The 2025 Annual Work Trend Index: The Frontier Firm is born · https://blogs.microsoft.com/blog/2025/04/23/the-2025-annual-work-trend-index-the-frontier-firm-is-born/ · Microsoft · 23-04-2025

Moderna: Leadership Team · https://investors.modernatx.com/leadership-team · Moderna, Inc.

AI-First Enterprise Operations: Reinventing the Operating System of Work · https://www.bcg.com/publications/2026/reinventing-the-operating-system-of-work-with-ai · Boston Consulting Group · 15-06-2026

When everyone has AI, your operating model will set you apart · https://www.accenture.com/en/i