DOCUMENT 01 · FOUNDATIONAL METHODOLOGY

SEVEN-G foundational methodology

Model, principles, governance cycles, decision gates and roles

Fernando García Varela Version 0.3 (working draft) 01-10-2026 Draft for review

10	2	8	6
principles	governance levels	phases per use case	roles with segregation of duties

IN THIS DOCUMENT

01 Purpose and scope

02 The SEVEN-G model

03 Principles

04 Basic concepts

05 Company level · Corporate cycle

06 Use case level · Lifecycle

07 Decision gates

08 Roles, bodies and segregation of duties

09 Application intensity: Express, Lite and Enterprise

10 Risks

11 Measurement

12 Nonconformities

13 Regulation and reference standards

14 How to declare that SEVEN-G is applied

15 Related documents

16 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regulation specific to each sector or jurisdiction may be incomplete, may not apply to a particular case or may become out of date as a result of regulatory changes, interpretations or supervisory positions after their consultation date.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 Always mandatory

The framework's normative reference. What no company may omit: the corporate cycle C1–C5, the eight phases with their gates and evidence (section 6), the gate decision rules (section 7), segregation of duties (section 8), the determination of Lite or Enterprise intensity for each use case (section 9) and the nonconformity process (section 12).

Obligation matrix and layered reading (document 94)

01

Purpose and scope

This document defines the SEVEN-G model and the rules that an organisation must apply in order to declare that it governs its artificial intelligence with this framework. It is the normative reference for all the other documents: the manuals, criteria, templates and tools develop what is established here and cannot contradict it.

1.1 What it contains

- The framework model: governance levels, components and the relationship between them.
- The principles and how they are verified.
- The corporate cycle, which governs AI in the company.
- The lifecycle of each use case, with its phases, evidence and decision gates.
- The decision rules, including the differences by ambition level.
- Roles, bodies and segregation of duties.
- The application intensity (Lite or Enterprise) and how it is determined.
- The integration of risk, measurement, nonconformities and regulation.

The operational detail of each topic is developed in the related documents (section 15).

1.2 Compatibility and modular adoption

SEVEN-G is compatible with governance structures that are already in place, with in-house or external consulting teams, and with pre-existing corporate frameworks for risk, control, technology or transformation. The framework does not require replacement; its function is to provide a common language, decision traceability and verifiable evidence to coordinate what the company already does with what it needs to strengthen.

Its design is modular: it can be implemented in full or by components. An organisation may start with one specific element (for example inventory, *gates,* measurement or board oversight) and progressively expand scope without breaking methodological coherence.

1.3 What is governed with SEVEN-G

Type of AI use	Examples	Treatment in SEVEN-G
AI use cases	Predictive models, generative AI solutions, agents, AI-enabled automations developed or adapted by the company.	Full lifecycle (phases 0–7) at the appropriate intensity.

Type of AI use	Examples	Treatment in SEVEN-G
Third-party AI embedded in processes	Supplier software with AI functions that play a part in decisions, operations or customer relationships.	Full lifecycle; the design and delivery phases focus on the selection, integration, contract and controls of the supplier.
Corporate use of general-purpose AI	AI-enabled assistants and productivity suites used by employees.	Inventory, acceptable use policy, training and technical controls. It moves to the full cycle if it meets any Enterprise criterion (section 9).
Unauthorised use	AI tools used without approval.	It is detected, recorded and regularised (authorisation, replacement or blocking) as a nonconformity.

1.4 Language conventions

Term	Meaning
must	Mandatory requirement. Failure to meet it is a nonconformity.
should	Recommendation. It may be omitted if the omission is justified and recorded.
may	Permitted option.

02

The SEVEN-G model

SEVEN-G stands for *Seven-phase Enterprise Value & Governance*. The name expresses the three ideas of the framework:

- **Seven value phases** that each use case goes through, preceded by a phase 0 that authorises it.
- **Enterprise value** as the criterion for all decisions.
- **Governance** present in every phase and in the company as a whole.

SEVEN-G architecture

Two connected levels of governance and four cross-cutting components

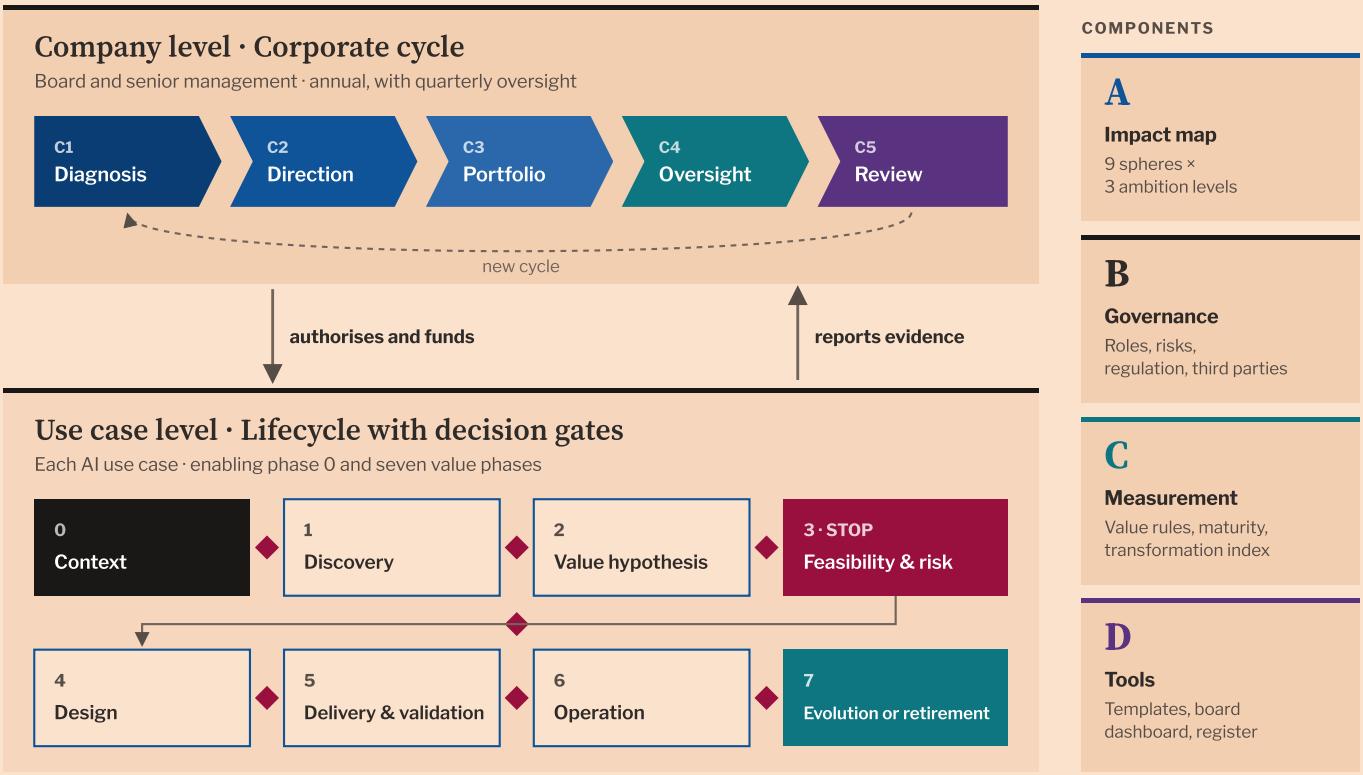

The diamonds represent the decision gates (*gates*). Source: SEVEN-G

The framework is organised into **two governance levels** and **four cross-cutting components**:

Element	What it is	Who uses it
Corporate cycle	Diagnosis, direction, portfolio, oversight and review of AI in the company.	Board, senior management, AI Committee.
Use case lifecycle	Phases 0–7 with auditable decision gates.	Use case teams, risk, AI Auditor, AI Committee.
A · Impact map	Nine spheres of impact and three ambition levels.	Board and management to decide where and how much to bet; teams to classify use cases.
B · Governance system	Roles, bodies, risks, regulation, third parties and nonconformities.	All levels.
C · Measurement system	Value rules, indicators, maturity and transformation index.	Board, AI Committee, management control, teams.
D · Tools	Templates, checklists, board dashboard and register of recommendations.	All levels.

The two cycles are connected: **the corporate cycle authorises and funds** use cases within a portfolio, and **each use case reports evidence** that feeds oversight and the annual review.

03

Principles

The principles guide decisions in situations that the rules do not cover. Each principle has a way of being verified.

#	Principle	What it means in practice	How it is verified
1	**Value before technology**	Every use case starts from a business problem or opportunity and is linked to a measurable outcome.	An approved value hypothesis exists before investing in construction.
2	**Governance is not optional**	Controls are designed in from the start and make it possible to move faster with less risk.	No use case under construction lacks an approved phase 0.
3	**Production is the only truth**	Value is demonstrated under real conditions, not in proofs of concept.	Value declared in production has a validation status.
4	**Reversibility**	Every solution can be stopped, rolled back or retired without compromising operations.	A tested rollback plan exists before go-live.
5	**Risk is systemic**	Risk is managed at portfolio level, as well as case by case.	The AI Committee reviews the concentration of risks in the portfolio.
6	**Evolve or retire**	No use case remains without evidence to justify it.	All use cases in production have a current continuity review.
7	**Segregation of duties**	Whoever builds does not control, and no one approves their own work.	*Gate* decisions record a decision-maker and a verifier who are not part of the team.
8	**Evidence, not declaration**	What is not documented and verified is not considered done.	Evidence has an author, date, version and verification.
9	**A conscious decision on ambition**	Efficiency and transformation are chosen, measured and decided separately.	Every use case has a classified and confirmed ambition level.
Sphere map and ambition levels (10)	**People at the centre of change**	AI is also assessed by its effect on work, capabilities and responsibilities.	Augment and Transform use cases have an adoption and people plan.

#	Principle	What it means in practice	How it is verified
Maturity model (11)	**Proportionality**	Governance effort matches each use case's risk, materiality and reversibility: no evidence is requested twice, no decision needs more participants than segregation of duties requires and no use case bears a disproportionate governance cost. What is never reduced is the core (section 14.1).	Each use case's intensity is determined and justified (section 9); evidence the company already has is referenced instead of copied (document 21, section 4); the cost of governance itself is measured and reviewed in C5 (document 41).

04

Basic concepts

The full definitions are in document 02 (Glossary).

Concept	Definition in SEVEN-G
AI system	A machine-based system that, with varying levels of autonomy, infers from the input it receives how to generate outputs —predictions, content, recommendations or decisions— that can influence physical or virtual environments. It is aligned with the definition in the EU AI Act.
Use case	An application of AI to a process, a product or a decision that pursues a value hypothesis through one or more AI systems. It is the unit that goes through the lifecycle and is identified with the code IA-AAAA-NNN.
Portfolio	The set of authorised use cases, with their budget, priority and balance of ambition.
Sphere	A domain of AI impact in the organisation. There are nine (document 10).
Ambition level	Type of bet: Optimise, Augment or Transform.
Intensity	The degree of rigour with which the cycle is applied: Lite or Enterprise.
Gate	A decision gate at the end of a phase, at which it is decided whether the use case proceeds and under what conditions.
Evidence	A verifiable document, record or result that demonstrates that a criterion has been met.
Dual validation	The rule whereby a *gate* requires both tangible results and verified documentation.
Nonconformity	Failure to meet a requirement of the framework, classified as minor, major or critical.

05

Company level · Corporate cycle

The corporate cycle is the mechanism through which the board and senior management exercise their responsibility for AI. It runs annually, with continuous oversight.

5.1 Stages

Stage	Objective	Mandatory outputs	Owner	Approves
Diagnosis (C1)	Understand the real situation of AI in the company.	AI system inventory; maturity assessment with evidence; current sphere map; transformation index profile; current validated value and cost.	AI Office	AI Committee; presented to the board
Direction (C2)	Decide where to play, with what ambition and with what risk appetite.	AI thesis; ambition level per sphere; risk appetite and thresholds (including those for Enterprise intensity and the return horizon); corporate AI policy; framework budget.	Senior management	Board
Portfolio (C3)	Select, prioritise and balance use cases.	Prioritised portfolio with ambition level, sphere, intensity, budget and owners; retirement criteria; available capacity.	AI Committee	AI Committee; the board approves Transform use cases
Oversight (C4)	Check that what was decided is being delivered and act on deviations.	Board dashboard; register of recommendations and decisions; outcomes of relevant *gates*; incidents and nonconformities; main portfolio risks.	AI Office	AI Committee (monthly); board or board committee (quarterly)
Review (C5)	Assess progress and adjust direction.	Review of maturity and the transformation index; delivery against the thesis; lessons learned; proposed adjustments for the new cycle.	AI Office with internal audit	Board

5.2 Governance calendar

Body	Frequency	What it reviews
Board of directors	Annual (C2, C5) and quarterly (C4)	Thesis, ambition, risk appetite, Transform use cases, board dashboard, transformation profile.
Board committee (audit, risk or technology)	Quarterly	Risks, regulatory compliance, incidents, major and critical nonconformities, audits.

Body	Frequency	What it reviews
AI Committee	Monthly	Portfolio, Enterprise *gates*, value and cost deviations, concentrated risks, retirements.
AI Office	Continuous	Inventory, methodology, consolidation of measurement, support for teams, preparation of information.
AI Auditor	At every Enterprise *gate* and by sampling in Lite	Verification of evidence and closure of nonconformities.

5.3 First implementation

An organisation adopting SEVEN-G for the first time should concentrate C1 to C3 into ninety days:

Period	Objective	Output
Month 1	Diagnosis	Initial inventory, maturity with evidence and transformation profile.
Month 2	Risks and opportunities	Main risks and opportunities by sphere, with owner, economic impact and timeframe.
Month 3	Governance structure	Bodies, roles, thresholds, *gates*, metrics and reporting cadence approved.

The detail is in document 90 (Implementation guide).

06

Use case level · Lifecycle

Each use case goes through eight phases. At the end of each one there is a decision gate (*gate*); in phase 6 the gate is replaced by a periodic continuity review.

Use case lifecycle

Eight phases with auditable decision gates and a real capacity to stop

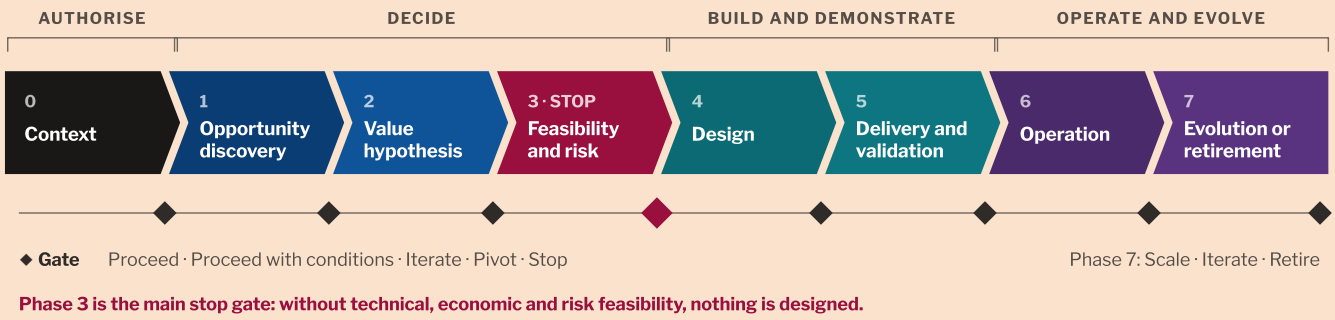

Source: SEVEN-G

6.1 Summary of phases and gates

Phase	Name	Question it answers	Decision gate
0	**Context and constraints**	Is the use case authorised, and within what framework?	Authorisation (G0)
1	**Opportunity discovery**	Is there a business opportunity that requires AI?	Opportunity (G1)
2	**Value hypothesis**	What value do we expect, how will we measure it and how will we know it has failed?	Hypothesis (G2)
3	**Feasibility and risk**	Is it technically, economically, regulatorily and organisationally feasible, with acceptable risk?	Feasibility (G3 · main stop gate)
4	**Solution design**	How is it built with control, human oversight and reversibility?	Design (G4)
5	**Delivery and validation**	Does it work and deliver value under real conditions?	Go-live (G5)
6	**Operation and governance**	Is it still working, delivering value and under control?	Continuity review (R6 · periodic)
7	**Evolution or retirement**	Do we scale, iterate or retire?	Scale or retire (G7)

6.2 Context and constraints (phase 0)

Objective	Formally authorise the use case and set the framework within which it will be developed.
Activities	Identify the strategic objective and the sphere; declare regulatory, ethical, data, budgetary and time constraints; assign the roles; determine the intensity (Lite or Enterprise); register the use case in the inventory.
Mandatory evidence	Use case charter · Context and constraints statement · Role assignment record · Intensity determination · Inventory registration.
Exit criterion	Committed sponsor, roles assigned without incompatibilities, fit with the AI thesis and the portfolio, and known constraints.
Specific rule	Without an approved G0, the use case is not authorised: it cannot consume budget or access production data.

6.3 Opportunity discovery (phase 1)

Objective	Identify opportunities from the business and discard those that do not require AI or have no plausible value.
Activities	Analyse the affected process or decision; identify non-AI alternatives; estimate the order of magnitude of the value; propose a sphere and ambition level using the criteria in document 12.
Mandatory evidence	Opportunity portfolio with screening notes · Non-AI alternatives considered · Proposed sphere and ambition level.
Exit criterion	The opportunity arises from a business need, AI contributes something that the alternatives do not, and the potential value justifies formulating a hypothesis.

6.4 Value hypothesis (phase 2)

Objective	Formulate a measurable and falsifiable value hypothesis. A hypothesis that cannot fail is not valid.
Activities	Define the primary and secondary metrics; measure the baseline; set the target and the success threshold; choose the attribution method (control group, before and after, or another justified method); estimate value in money with a formula; confirm the ambition level; define the stop criteria.
Mandatory evidence	Value hypothesis canvas · Baseline metrics · Attribution method · Ambition level confirmation · Stop criteria.

Exit criterion	Falsifiable hypothesis, measured baseline (not estimated unless justified), value expressed in money with a formula, and stop criteria defined before investing.
Specific rule	In Transform use cases, the hypothesis may carry greater uncertainty, but must include learning milestones, an investment limit per stage and an explicit decision by the board.

6.5 Feasibility and risk (phase 3)

Objective	Decide whether the use case is feasible with acceptable risk. This is the main stop gate of the cycle.
Activities	Assess technical feasibility and the availability and quality of the data; estimate full costs (build, recurring and adoption); classify the use case under the applicable regulation; carry out the relevant impact assessments (data protection, fundamental rights); identify and assess risks; assess suppliers; assess the impact on people.
Mandatory evidence	Feasibility assessment · Regulatory classification · Applicable impact assessments · Risk matrix and register · Mitigation and contingency plan · Supplier assessment, where there are suppliers.
Exit criterion	Feasibility demonstrated with real data; no critical risk without accepted mitigation; regulatory classification carried out with legal judgement; expected net value consistent with the risk appetite and the horizon set in C2.
Specific rule	Practices prohibited by regulation do not proceed beyond this phase under any circumstances.

6.6 Solution design (phase 4)

Objective	Design a solution that is controllable, subject to oversight and reversible.
Activities	Define the architecture; document data and model lineage; design human oversight (what the system decides, what a person validates and what is never delegated); design security controls, including agent identity, permissions and limits of action; define monitoring; prepare the rollback plan; design adoption.
Mandatory evidence	Architecture record · Data and model lineage · Governance and human oversight design · Security design · Rollback plan · Adoption plan.
Exit criterion	The design covers the controls required by the risk classification; human oversight is defined; a stop mechanism exists; the phase 3 risks have a designed control.

6.7 Delivery and validation (phase 5)

Objective	Build, test and demonstrate value under real conditions before final go-live.
Activities	Build or integrate the solution; carry out functional, performance, bias, robustness and security testing (including prompt injection testing for generative AI and agents); run a pilot with measurement according to the attribution method; test the rollback plan; train users.
Mandatory evidence	Delivery report · Validation and test results · Pilot results against the hypothesis · Rollback plan test · Updated risk register · Go-live sign-off record.
Exit criterion	The pilot results meet the success threshold, or do so with accepted conditions; critical controls work; operations are ready.
Specific rule	The go-live of an Enterprise use case requires **multi-level sign-off with veto power**: technical owner, risk and compliance, information security and data protection. Each sign-off is recorded.

6.8 Operation and governance (phase 6)

Objective	Operate stably, maintain control and keep measuring value.
Activities	Monitor performance, degradation, bias, costs and security; manage incidents; carry out the post-market monitoring required by regulation; measure realised value; manage changes.
Mandatory evidence	Operations manual · Monitoring and alerting configuration · Incident response plan · Incident and change log · Value realisation tracking.
Continuity review (R6)	At least quarterly in Enterprise and half-yearly in Lite. It checks realised value against the hypothesis, stability, incidents, compliance and whether the risk classification is still valid. If it detects significant deviations, it brings forward the G7 *gate*.

6.9 Evolution or retirement (phase 7)

Objective	Decide, on the basis of evidence, whether the use case is scaled, iterated or retired.
Activities	Consolidate realised value and its validation status; review the actual ambition level against the declared one; assess accumulated risks; if it is retired, plan the retirement.
Mandatory evidence	Value realisation tracking · Scaling or retirement decision record · Lessons learned · Retirement plan, where applicable.

Possible outcomes	**Scale** (new phase 0 for the extended scope), **Iterate** (return to the relevant phase; if kept unchanged, return to phase 6) or **Retire**.
Specific rule	Every retirement records the date, reason, deciding body, replacement if any, treatment of data and models, and communication to those affected.

6.10 Evidence map

Phase	Mandatory evidence
0	Use case charter · Context and constraints statement · Role assignment record · Intensity determination · Inventory registration
1	Opportunity portfolio and screening notes · Non-AI alternatives · Proposed sphere and ambition level
2	Value hypothesis canvas · Baseline · Attribution method · Ambition confirmation · Stop criteria
3	Feasibility assessment · Regulatory classification · Impact assessments · Risk matrix and register · Mitigation and contingency plan · Supplier assessment
4	Architecture record · Data and model lineage · Governance and human oversight design · Security design · Rollback plan · Adoption plan
5	Delivery report · Validation and test results · Pilot results · Rollback test · Updated risk register · Go-live sign-off
6	Operations manual · Monitoring and alerting · Incident response plan · Incident and change log · Value tracking
7	Value realisation tracking · Scaling or retirement decision · Lessons learned · Retirement plan

The templates for each piece of evidence form part of block H of the documentation library.

6.11 Lifecycle traceability

The AI portfolio is managed as a funnel, in the same way that a sales organisation manages its opportunities. To this end, every use case must be registered in the **use case register** (document 03, tool T01), and the register must retain, as a minimum:

- The **entry and exit date of each phase**, and the periods on hold with their reason.
- The **request, verification and decision for each *gate***, with its outcome, iteration, decision-maker and verifier.
- The **status of each *gate* criterion** (met, not met, not applicable or pending) with the linked evidence.

- The **conditions** imposed, with deadline, owner and status.
- The **controlled taxonomy tags**: sphere, ambition level, intensity, regulatory classification, technology, exposure and value type.
- The **coded reason** for every stop or retirement.

In C2 the company approves the **reference time limits per phase** and the maximum decision time for a *gate*. Use cases that exceed them are flagged as stalled and reviewed by the AI Committee.

07

Decision gates

7.1 How a gate works

Each *gate* follows the same sequence: the team provides the evidence, the verifier checks that it is complete and valid, and the competent body decides.

Logic of a decision gate

Decisions are only taken on verified evidence

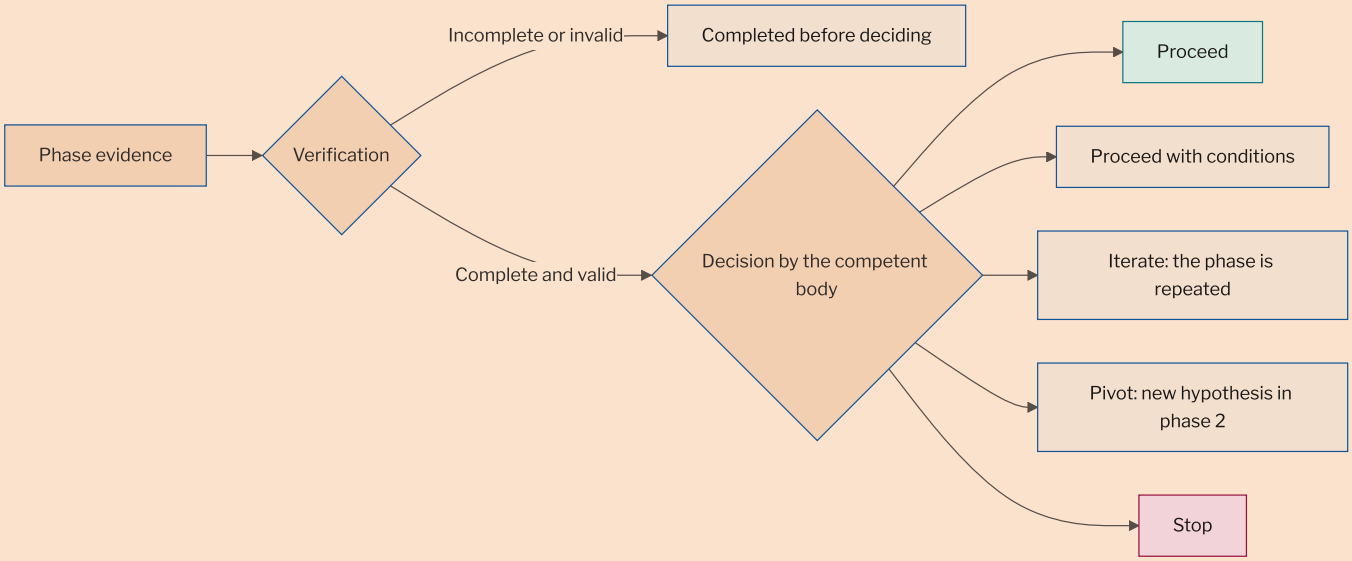

Source: SEVEN-G

7.2 Dual validation

A *gate* can only be passed if two conditions are met at the same time:

1. **Tangible results**: the phase criteria are met with real data, tests or results.

2. **Verified documentation**: the mandatory evidence exists, is traceable (author, date and version) and has been verified by the appropriate person.

Results without documentation, or documentation without results, do not pass the *gate*.

7.3 Possible outcomes

Outcome	When it applies	Consequence
Proceed	All criteria are met and all evidence is verified.	The use case moves on to the next phase.
Proceed with conditions	The essential criteria are met and non-critical aspects remain pending.	It moves on to the next phase with explicit conditions, a deadline and an owner. The conditions are verified at the next *gate*. Not permitted for critical security, legal compliance or human oversight controls.
Iterate	The results are insufficient but the hypothesis remains plausible.	The phase activities are repeated and the use case returns to the same *gate*.
Pivot	The hypothesis does not hold, but a reasonable alternative exists. Only at G1, G2 and G3.	The use case returns to phase 2 with a new hypothesis, keeping the approved context.
Stop	There is no plausible value, feasibility is not demonstrated or the risk is unacceptable.	The use case is closed, lessons learned are recorded and resources are released.
Scale	Only at G7: value demonstrated and risk under control.	New phase 0 for the extended scope.
Retire	Only at G7: the value is not sustained, the risk has increased or a better alternative exists.	Planned and recorded retirement.
Proceed with operation · Proceed with conditions · Bring G7 forward	Only at R6: the use case keeps delivering value under control, does so with correctable deviations, or shows relevant deviations in value, risk or compliance.	Remains in production; with conditions, deadline and owner; or G7 is opened without waiting for the schedule.

A well-founded decision to **stop** is a valid outcome of the method. Avoiding a bad investment also creates value.

> **In plain words: what does "iterate" mean?** Iterating means repeating the work of the phase the use case is currently in, without moving on to the next one, because the result does not yet meet the required criterion but the underlying idea still looks valid. It is not starting from scratch (that would be pivoting, with a different hypothesis) and it is not closing the use case (that would be stopping): it is correcting and trying again at the same point. For example, in phase 2 a value hypothesis is tested and the measured saving falls short of the target, but the team identifies a concrete adjustment (a different user segment, a different process) that could close the gap; instead of moving on to phase 3 with insufficient data, the use case iterates: it repeats phase 2 with that adjustment and is presented again to the same *gate*. **Why only twice is allowed (rule 5 in section 7.4):** every iteration has a cost in time and budget, and if two attempts to fix the same thing are not enough, that is a sign the problem is not a minor adjustment but something more fundamental (the hypothesis, the design or the sponsorship). That is why, the third time, the decision is no longer made by whoever runs the use case: it **is escalated to whoever has the authority to decide with more context** (for example, from the sponsor to the AI committee), who assesses whether it is worth iterating again, pivoting to another hypothesis, or stopping.

7.4 Decision rules

1. **No one decides on their own work.** Whoever provides evidence neither verifies it nor decides on it.
2. **No mandatory evidence, no decision.** The competent body cannot decide on incomplete evidence.
3. **Evidence must exist before the *gate*.** Documentation produced after the fact to justify progress already made invalidates the *gate* and constitutes a major nonconformity.

4. **Conditions have a deadline and an owner.** A condition that expires unmet turns the outcome into **Iterate**.
5. **Iteration limit.** After two iterations at the same *gate*, the decision is escalated to the higher body.
6. **Stop criteria are set in advance.** They cannot be relaxed during the phase to avoid a decision to stop without the approval of the body that authorised the use case.
7. **Every decision is recorded** in the *gate* decision record, with outcome, reason, conditions, decision-maker and verifier.

7.5 Who verifies and who decides

Gate	Verifies (Lite)	Decides (Lite)	Verifies (Enterprise)	Decides (Enterprise)
Authorisation (G0)	AI Office	Sponsor	AI Auditor	AI Committee
Opportunity (G1)	AI Office	Sponsor	AI Auditor	Sponsor, informing the committee
Hypothesis (G2)	AI Office	Sponsor	AI Auditor	AI Committee
Feasibility (G3)	AI Office	Sponsor with risk clearance	AI Auditor	AI Committee
Design (G4)	AI Office	Sponsor with risk clearance	AI Auditor	AI Committee
Go-live (G5)	AI Office	Sponsor with risk clearance	AI Auditor	AI Committee after multi-level sign-off
Continuity (R6)	AI Office	Sponsor	AI Auditor	AI Committee
Scale or retire (G7)	AI Office	Sponsor	AI Auditor	AI Committee

Transform use cases additionally require board approval at G2 (authorisation of the bet) and at G7 when the decision is to scale. At **Lite** intensity, G0, G1 and G2 may be resolved in a single session, as may G4 and G5, provided that each piece of evidence is verified. Under **Lite**, the AI Office verifies the gate; the risk, regulatory classification and supplier evidence (P10 to P14) and the change and validation evidence (P21 and P22) are verified by the AI Risk Owner, unless the AI Risk Owner prepared them: then the AI Office or the AI Auditor verifies them.

7.6 Criteria differentiated by ambition level

Applying the same return criteria to an efficiency use case and to a transformation bet blocks the latter or inflates the former. That is why the key *gates* are assessed differently:

Gate	Optimise	Augment	Transform
Hypothesis (G2)	Baseline for cost, time or errors; expected savings with a formula.	Performance and cost metrics; adoption target.	Return hypothesis with learning milestones; investment limit per stage; board approval.
Feasibility (G3)	Positive expected annual net value and NPV ≥ 0 within the horizon and at the rate set in C2 (document 40).	Feasibility of adoption and of the role change, in addition to the expected net value.	Feasibility of the first stage; stop criteria per stage; documented option value.
Go-live (G5)	Efficiency validated against the baseline; plan to realise the released capacity.	Actual adoption and performance improvement measured.	Market or customer evidence: usage, conversion, initial revenue or verified operational change.
Scale or retire (G7)	Realised savings, not just released capacity.	Sustained performance and reassigned capacity.	Measured return and verified change in the operating model or the offering.

08

Roles, bodies and segregation of duties

8.1 Use case roles

Role	Function	Responsibility	Cannot
AI Sponsor	Decides	Accountable for value and investment; champions the use case before the decision-making bodies.	Verify evidence or act as risk owner or auditor of their own use case.
AI Product Owner	Builds	Accountable for the value hypothesis, actual use and adoption.	Verify evidence or decide *gates* for their own use case.
AI Technical Owner	Builds	Accountable for the solution, the data, the models and their technical documentation.	Verify evidence or decide *gates* for their own use case.
AI Operations Owner	Builds	Accountable for stability, monitoring, incidents and changes in production.	Verify evidence of their own operations.

Role	Function	Responsibility	Cannot
AI Risk Owner	Controls	Accountable for the assessment and monitoring of risk and compliance, and issues clearance.	Be part of the team that builds the use case.
AI Auditor	Controls	Verifies evidence at Enterprise *gates* and closes nonconformities.	Take part in the design, construction or operation of the use case, or report hierarchically to the sponsor.

One person may take on more than one role if the roles are not incompatible (section 8.2). In small organisations, the AI Auditor may be external or come from internal audit.

8.2 Incompatibilities

	Sponsor	Product	Technical	Operations	Risk	Auditor
Sponsor	—	Compatible in Lite	Compatible in Lite	Compatible in Lite	Incompatible	Incompatible
Product		—	Compatible	Compatible	Incompatible	Incompatible
Technical			—	Compatible	Incompatible	Incompatible
Operations				—	Incompatible	Incompatible
Risk					—	Incompatible in Enterprise
Auditor						—

8.3 Company bodies

Body	Indicative composition	Functions in SEVEN-G
Board of directors or board committee	Directors; may be supported by a director or adviser with AI experience.	Approves the AI thesis, the ambition, the risk appetite and Transform use cases; exercises quarterly oversight; reviews progress annually.
AI Committee	Senior management from business, technology, data, risk, compliance, security, data protection and people.	Manages the portfolio; decides Enterprise *gates*; reviews portfolio risks; decides retirements; escalates to the board as appropriate.

Body	Indicative composition	Functions in SEVEN-G
AI Office	Small methodology, portfolio and measurement team.	Maintains the inventory and the methodology; verifies the gate criteria in Lite-intensity use cases (in Enterprise ones the AI Auditor verifies them); consolidates measurement; prepares information for the bodies.
Second line	Risk, compliance, information security, data protection.	Provides the risk owners; issues clearances; signs off Enterprise go-live.
Third line	Internal audit or external auditor.	Provides or supervises the AI Auditors; audits compliance with the framework.

SEVEN-G does not create a parallel governance structure: the bodies may be existing committees with an extended mandate.

8.4 Responsibilities by phase

A is accountable for the outcome · **R** does the work · **C** is consulted · **I** is informed · **V** verifies the evidence at the *gate*.

Phase	Sponsor	Product	Technical	Operations	Risk	Auditor
0 · Context	A	R	C	I	C	V
1 · Discovery	A	R	C	I	C	V
2 · Value hypothesis	A	R	C	I	C	V
3 · Feasibility and risk	A	R	R	C	R	V
4 · Design	I	C	A/R	C	C	V
5 · Delivery and validation	I	A	R	C	C	V
6 · Operation	I	C	C	A/R	C	V
7 · Evolution or retirement	A	R	C	C	C	V

09

Application intensity: Express, Lite and Enterprise

9.1 How it is determined

Intensity is determined in phase 0 and reviewed at G3 and at every continuity review. Meeting a single Enterprise criterion is enough for that intensity to apply. With no Enterprise criterion, the use case is Lite or, if it also meets the conditions in section 9.4 and the risk function confirms it, Express.

Determining the intensity

Assessed in phase 0 and reviewed at G3 and at each continuity review. A single “yes” is enough to apply Enterprise; Express only if, in addition, its eligibility is met (01 §9.4)

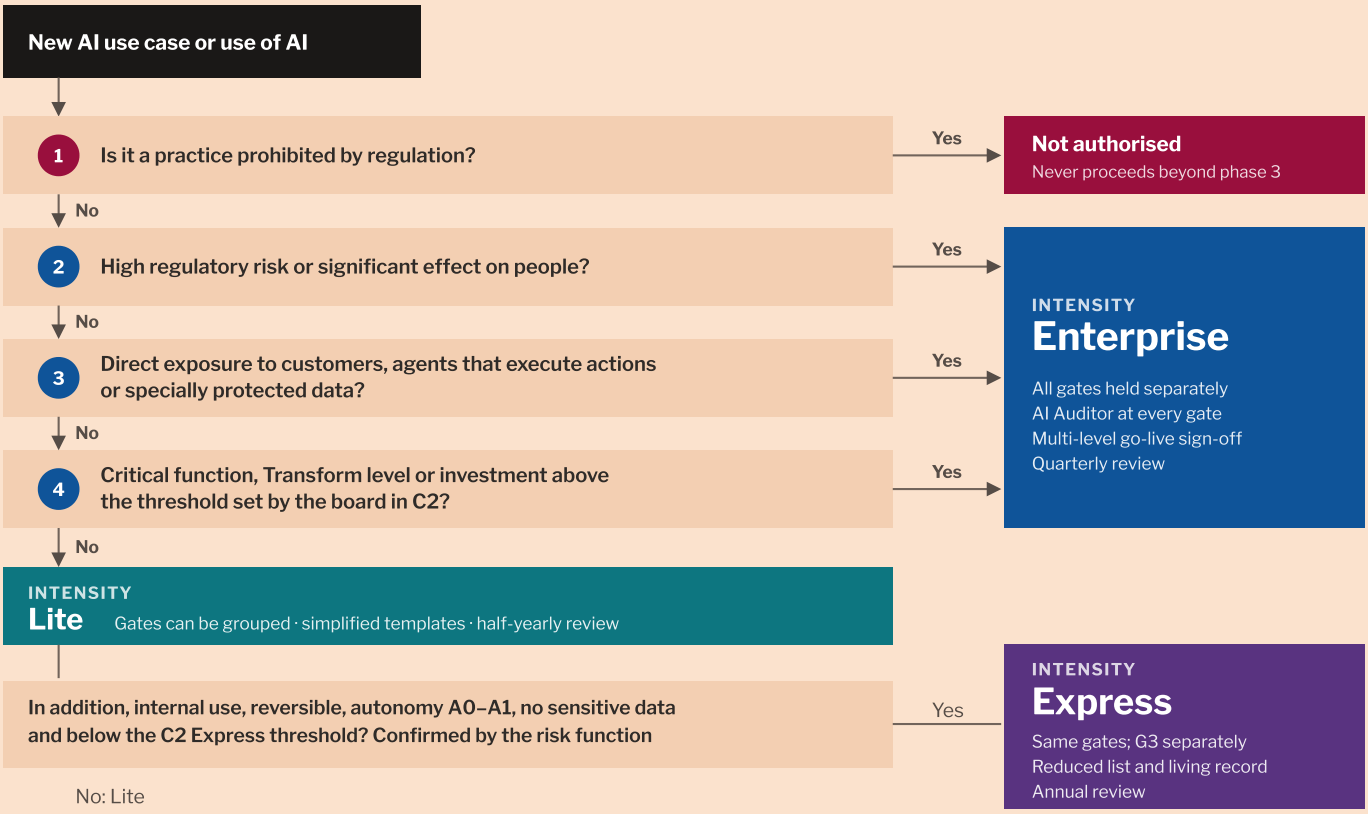

Source: SEVEN-G

9.2 Enterprise criteria

Criterion	Description
High regulatory risk	The system is classified as high-risk under the applicable regulation.

Criterion	Description
Decisions about people	The output significantly influences decisions that affect people (employment, credit, insurance, access to services, among others).
Direct exposure	Customers, patients, citizens or other external persons interact directly with the system.
Agents with the ability to act	The system executes actions —writing to systems, sending communications, making payments— without a person validating each action before it is executed (autonomy A2 or A3, document 35), and those actions affect third parties, money, personal data or production systems.
Specially protected data	It processes special categories of personal data or critical confidential information.
Critical function	It supports a critical or important business function, or one subject to sector-specific operational resilience regulation.
Transform level	The use case is classified as Transform.
Investment	The investment exceeds the threshold approved by the board in C2.

9.3 What changes between intensities

Aspect	Express	Lite	Enterprise
Criteria	Reduced list (21 §2.5): all applicable "Yes ◆" criteria and those that sustain the core	Lite column of document 21	All
Gates	Same gates as Lite: G0–G2 and G4–G5 in one decision each; G3 separately, may be resolved in writing	G0–G2 and G4–G5 may be grouped	All separately
Evidence	Living record (P01, Express section) and referenced company evidence (21 §4.3)	Simplified templates	Full templates
Verification	AI Office	AI Office; AI Auditor by sampling	AI Auditor at all *gates*
Decision	Sponsor, with risk clearance at G3, G4 and G5	Sponsor, with risk clearance at G3, G4 and G5	AI Committee; board for Transform
Go-live	Risk clearance	Risk clearance	Multi-level sign-off with veto
Continuity review	Annual	Half-yearly	Quarterly

Aspect	Express	Lite	Enterprise
Visibility	Aggregated board dashboard	Aggregated board dashboard	Board dashboard by use case

9.4 Express intensity

Express intensity exists so that a small, internal and reversible use case does not bear the same administrative path as one of hundreds of hours (principle 11, proportionality). **It changes what is assessed at each gate, not which gates there are**: the use case goes through all the gates with their record (N-04) and G3 is still decided separately (N-05).

Eligibility. No Enterprise criterion in section 9.2 and, in addition, all these conditions: internal use (no direct exposure to customers or external people); reversible within days, with no residual effect; autonomy A0 or A1; no special categories of data or critical confidential information; investment plus annual cost below the **Express threshold** approved in C2; and, if an AI supplier is involved, that it is already approved (current P55). **It is confirmed by the risk function** (in small companies, the AI Office), never by the sponsor alone. It is recorded in P04 and in T04.

Path. Three decisions with the same gates: *entry* (G0, G1 and G2 in one decision), **G3 separately**, which at Express may be resolved in writing —the sponsor decides with the risk clearance recorded in T01—, and *go-live* (G4 and G5 in one decision). Each gate is assessed with the **reduced list of Express criteria** in document 21 (section 2.5): all applicable "Yes ◆" criteria and those that sustain the core. The evidence is a **living record** (P01, Express section) plus the company evidence that is referenced (21 §4.3). The continuity review is **annual** and may bring G7 forward.

Automatic exit. If any Enterprise criterion appears or an eligibility condition ceases to be met, the use case moves to Lite or Enterprise from that moment, with a classification change event, and the following gates are assessed with the list of its new intensity. The same applies if the annual review does not demonstrate value.

What does not change. The whole core (section 14.1): inventory and register, segregation of duties, dual validation with evidence prior to the decision, "Yes ◆" criteria, value hypothesis with baseline and stop criterion, classification and risks, tested rollback.

> **Why it matters.** With only Lite and Enterprise, a low-cost internal assistant goes through 92% of the criteria of an agent that serves customers. That disproportion pushes people not to register small cases, which is exactly what governance wants to avoid. Express keeps traceability and critical controls and removes the rest, and its automatic exit prevents it from becoming a shortcut.

10

Risks

SEVEN-G manages risks at two levels: in each use case (phase 3 and continuous monitoring) and in the portfolio (AI Committee). The full methodology is in document 33.

Category	Examples
Strategic	Misalignment with the AI thesis; dependence on a single bet; missed opportunity.
Technical	Insufficient performance; degradation; lack of robustness; hallucinations in generative AI.
Data	Quality, availability, bias, legal basis for use.
Economic	Cost overruns; unrealised value; rising recurring costs.
Legal and compliance	Incorrect regulatory classification; failure to meet transparency, oversight or data protection obligations.
Organisational	Lack of adoption; loss of knowledge; unmanaged effect on people.
Reputational	Discriminatory outcomes; errors visible to customers; use perceived as inappropriate.

These categories are complemented by three areas that require specific controls:

- **Generative AI and agents**: prompt injection, information leakage, excessive permissions, unauthorised actions, lack of traceability of the intent behind each action.
- **Exposure to attacks that use AI**: impersonation, attack automation, accelerated exploitation of vulnerabilities. The detail is in document 35.
- **AI third parties and suppliers**: technological dependency, use of data by the supplier, uncommunicated model changes, concentration and exit. The detail is in document 36.

11

Measurement

Measurement is integrated into the lifecycle and the corporate cycle. The value measurement rules (document 00, section 6, and document 40) apply in all cases.

Moment	What is measured	Purpose
Phase 2	Baseline, target and success threshold; expected value in money with a formula.	Make the hypothesis falsifiable.
Phase 3	Full costs and expected net value.	Decide on economic feasibility.
Phase 5	Pilot results using the attribution method.	Demonstrate value before production.
Phase 6	Realised value, actual costs, performance, incidents, adoption.	Maintain control and detect deviations.
Phase 7	Validated realised value; actual versus declared ambition.	Decide on scaling or retirement.
C1 and C5	Maturity with evidence; transformation index; portfolio value and cost.	Diagnose and review direction.
Oversight (C4)	Board dashboard; decision agility (time from idea to approval and to production).	Oversee.
Continuous	Funnel metrics: time in phase, decision time, conversion per *gate*, iterations, stalled use cases, expired conditions and stop reasons (document 03).	Manage the portfolio as a funnel and detect bottlenecks.

Every declared amount has one of three statuses: **validated** (by management control or audit), **declared** (by the responsible area) or **estimated** (by the committee, the board or the assessment team).

12

Nonconformities

A nonconformity is any failure to meet a mandatory requirement of the framework. It is managed through a single process; the detail is in document 37.

Nonconformity management

A single process for all deviations. Only the AI Auditor can close a nonconformity

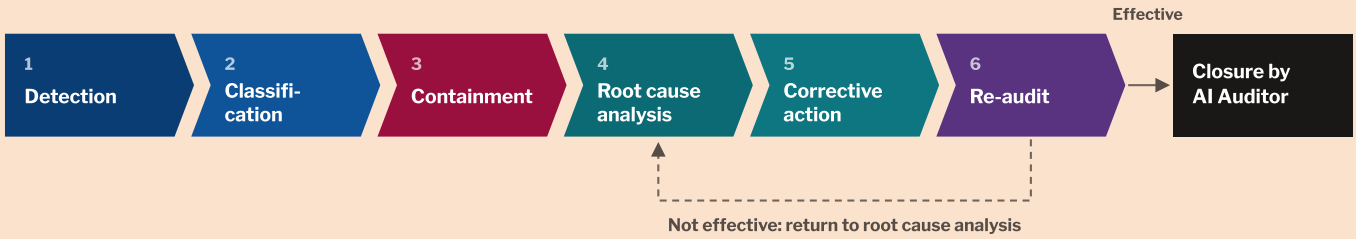

Source: SEVEN-G

Type	Examples	Containment	Action plan	Reports to
Critical	System in production without an approved *gate*; prohibited practice; unreported serious incident; critical control disabled.	Immediate (within 48 hours at most), including shutting down the system if necessary.	Within 10 days at most.	AI Committee and board committee.
Major	Evidence produced after the fact; self-approval; expired condition on a relevant control; omitted continuity review.	Within 10 days at most.	Within 30 days at most.	AI Committee.
Minor	Incomplete evidence with no impact on the decision; delay in updating records.	Not required.	Before the next *gate* or review.	AI Office.

The time limits are for reference and the company may adjust them in C2, without exceeding those set by the applicable regulation.

13

Regulation and reference standards

SEVEN-G **maps** regulation instead of embedding it: each obligation is linked to a phase, a role and a piece of evidence. In this way, when a regulation changes, the mapping is updated without rebuilding the framework. The detailed mapping is in document 34.

Reference	Main fit in SEVEN-G
EU AI Act	Classification by risk level and prohibited practices (phase 3); risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy and robustness (phases 3 to 6); AI literacy of staff (C2 and people); post-market monitoring and reporting of serious incidents (phase 6).
ISO/IEC 42001 (AI management system)	The corporate cycle acts as the management system: context and leadership (C1–C2), planning (C2–C3), operation (lifecycle), performance evaluation (C4–C5) and improvement (C5 and nonconformities).
NIST AI RMF	Govern (corporate cycle and roles), map (phases 1–3), measure (phases 2, 5 and 6) and manage (phases 3–7).
NIST CSF 2.0 and **Cyber AI Profile** (the latter in draft)	The security of AI systems is expressed with the six CSF functions (govern, identify, protect, detect, respond and recover): SEG and AG controls of document 35 in phases 3 to 6 and exposure to offensive AI in C4. The Cyber AI Profile is used for guidance while it remains a draft and does not underpin *gate* criteria (34 §5.3).
GDPR	Legal basis and data minimisation (phase 3), data protection impact assessment (phases 3–4), data subject rights and automated decisions (phases 4–6).
DORA and **NIS2**, where applicable	Technology supplier risk (phases 3–4), incident management and reporting (phase 6), resilience of critical functions (Enterprise criterion).
Sector-specific regulation	Incorporated into the context statement (phase 0) and the company's regulatory mapping.

This document does not constitute legal advice. The regulatory classification of each system must be carried out with qualified legal judgement.

14

How to declare that SEVEN-G is applied

14.1 The SEVEN-G core

The **core** is the minimum, stable set of rules that distinguishes applying SEVEN-G from using some of its templates. It is the same for every company, intensity and scope. Everything else in the library is either a **practice** —the recommended way, or the mandatory way when its trigger occurs, of meeting the core (document 94, section 2.3)— or an **instrument** —a template or a tool that makes it easier—. The core has its own version and only changes by express decision of the framework's author, which is recorded; documents, templates and tools may evolve without changing it.

Code	Core rule	Where it is developed
N-01	**Complete inventory** of AI systems —own, third-party, corporate use and unauthorised— with an owner, regulatory classification and intensity.	document 32; document 90, section 2.4
N-02	**Every use case is registered** before consuming budget, with its dates per phase, decisions, criteria and closing reason (traceability in section 6.11).	section 6.11; document 03
N-03	**Intensity is determined for each use case** in phase 0 and reviewed at G3 and at each R6. One Enterprise criterion is enough.	section 9
N-04	**No phase or gate is skipped**: every use case goes through all the gates of its lifecycle and each one is recorded with its criteria and its outcome. Which criteria are assessed and which sessions may be grouped depends on the intensity.	sections 7 and 9; document 21, section 3.4
N-05	**G3 is always decided separately**: it is the main stop gate.	section 6.1; document 90, section 2.4
N-06	**Segregation of duties**: roles and bodies assigned in line with the incompatibilities in section 8.2; whoever builds neither verifies nor decides their own work, and risk clearance is independent of the team.	section 8; document 30
N-07	**Dual validation and evidence prior to the decision**: without verified evidence there is no decision.	section 7.2; document 21, sections 3 and 4
N-08	**The "Yes ◆" criteria** (security, legal compliance, human oversight) are never accepted as a condition.	document 21, section 2.2
N-09	**Value hypothesis with baseline and stop criteria** before investing in feasibility.	section 6.4; document 40
N-10	**Regulatory classification and risk register** before design; a Critical residual risk without board approval blocks G3 and G5.	document 32; document 33
N-11	**Tested rollback plan and current continuity review** for everything in production; whatever was already in production when the framework was adopted is regularised within the time limit approved in C2, with a review equivalent to G7.	section 6.8; document 52; document 14, section 11
N-12	**The board approves in C2**, after the C1 diagnosis, the AI thesis, the ambition per sphere and the risk appetite; it receives AI information at least quarterly and approves Transform bets.	section 5; document 13
N-13	**Value and cost figures follow the measurement rules** and are reported to the board with the oversight dashboard.	section 11; document 40; document 60
N-14	**Nonconformities are managed** with the process in section 12.	section 12; document 37

Core version: **1 (01-10-2026)**. Document 94, section 3, repeats this list for convenience; if they differ, this one prevails.

> **Why it matters.** A framework that grows without a stable core forces people to reread it in full to know what is non-negotiable, and every new document looks like one more obligation. With a small core, a company knows what must never be missing and can apply the rest in proportion to the risk and materiality of each use case (principle 11), without the library having to stop growing as guidance.

14.2 Declaration

An organisation may declare that it applies SEVEN-G when it meets the fourteen core rules (N-01 to N-14). The declaration is drafted with P61, cites the core version it applies and is audited in accordance with 38 §11.

15

Related documents

Document	Relationship to this document
00 · What SEVEN-G is and how it helps companies	Framework overview and value measurement rules.
02 · Glossary	Full definitions of the terms.
03 · Tools and use case register	Use case register, funnel metrics, data model and tool catalogue.
10 · Sphere map and ambition levels	Component A.
11 · Maturity model	Diagnosis in C1 and C5.
12 · Transformation index	Ambition classification and company profile.
13 · AI thesis and risk appetite	C2 outputs.
14 · Portfolio management	C3 outputs and retirement criteria.
20 · Phase manuals	Operational development of phases 0–7.

Document	Relationship to this document
21 · *Gate* and audit criteria	Detailed criteria by gate, intensity and ambition.
22 · Checklists	Binary controls per *gate*.
30 · Governance model	Development of roles, bodies and escalation.
33 · AI risk methodology	Matrix, register and typical risks.
34 · Regulatory mapping	Obligation by obligation.
35 · AI and agent security	Agent-specific controls and exposure to attacks.
37 · Nonconformities and incidents	Development of the process in section 12.
40 · Value measurement rules	Development of the measurement system.
90 · Implementation guide	First ninety days and regularisation.

16

Version control

Version	Date	Changes
0.1	16-09-2026	First version. Consolidates the previous foundational methodology and the *gate* criteria into a single model; incorporates the corporate cycle, the criteria differentiated by ambition level, intensity determination, multi-level go-live sign-off, the incompatibility table, nonconformity time limits, regulatory fit and lifecycle traceability as a funnel.
0.2	01-10-2026	Principle 11, proportionality (section 3). Section 14 defines the SEVEN-G core (14.1, rules N-01 to N-14, version 1), which unifies the seven declaration conditions and the twelve rules never omitted in document 94, and the declaration (14.2).
0.3	01-10-2026	Express intensity (section 9.4, D147): eligibility confirmed by risk, same gates with G3 separately, reduced list of criteria, living record, annual review and automatic exit; Express column in 9.3 and in the intensity figure.

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) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) · https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng · Parlamento Europeo y Consejo de la UE · 27-04-2016

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

Cybersecurity Framework Profile for Artificial Intelligence (Cyber AI Profile): NIST Community Profile, NIST IR 8596 (Initial Preliminary Draft) · https://csrc.nist.gov/pubs/ir/8596/iprd · National Institute of Standards and Technology (NIST), NCCoE · 16-12-2025

ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system · https://www.iso.org/standard/42001 · ISO/IEC (JTC 1/SC 42) · 18-12-2023

Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 · https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10 · National Institute of Standards and Technology (NIST), Departamento de Comercio de EE. UU. · 26-01-2023

The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29 · https://csrc.nist.gov/pubs/cswp/29/the-nist-cybersecurity-framework-csf-20/final · National Institute of Standards and Technology (NIST), Departamento de Comercio de EE. UU. · 26-02-2024