DOCUMENT 61 · BOARD CONVERSATION GUIDE

Board conversation guide

How to talk about AI with a board of directors: language, response format, questions and session script

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

4	3	9	60
response formats	use cases in detail per session	spheres with questions for the board	minutes of session script

IN THIS DOCUMENT

01 Purpose and scope

02 Principles of the conversation

03 Business language

04 Response format for the board

05 How many use cases to present

06 Questions by sphere

07 Questions by stage of the corporate cycle

08 Script for a 60-minute session

09 Frequent difficult questions

10 Role of the director or adviser with AI experience

11 Frequent mistakes

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regulation specific to each sector or jurisdiction may be incomplete, may not apply to a particular case or may become out of date as a result of regulatory changes, interpretations or supervisory positions after their consultation date.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

Guide for preparing and conducting the conversation with the board: business language, response format, questions per sphere and session script. It creates no obligations.

Obligation matrix and layered reading (document 94)

01

Purpose and scope

This guide helps whoever presents AI to a board of directors or its board committee —senior management, the chair of the AI Committee, the AI Office, a director or adviser with AI experience— to hold a conversation that is **useful for deciding and overseeing**.

It contains:

- The language to use and the terms best avoided.
- The response format for the board (common specification §5.8), with examples.
- How many use cases to present and how to choose them.
- Reference questions by sphere (all nine) and by stage of the corporate cycle (C1–C5).
- A detailed script for a 60-minute session and its variant for oversight sessions.
- The most frequent difficult questions and how to answer them with evidence.
- The role of the director or adviser with AI experience.
- The most frequent mistakes.

The content of the pack delivered before the session is in document 60.

02

Principles of the conversation

#	Principle	What it means in the room
1	**The board decides on business, not on technology**	The discussion is about risk, opportunity, accountability, competitive advantage and money. Technology only appears when it changes one of those things.
2	**The answer comes first**	Every answer starts with one of the four formats in section 4. Then comes the reason and the evidence.
3	**Every figure has a status**	Validated, declared or estimated. If it is not known, say "no data" and when the data will be available.
4	**Few use cases, well told**	Three use cases in detail at most. The rest are in the dashboard.
5	**Say what is not working**	Stops, retirements, incidents and unmaterialised value are presented with the same clarity as successes.
6	**Every board question ends up in the register**	If a question cannot be answered in the session, it becomes a recommendation or an assignment with an owner and a date (document 62).

#	Principle	What it means in the room
7	**Nobody marks their own exam**	Whoever presents a use case does not assess its compliance; the assessment is made by someone independent.

03

Business language

3.1 Rules

1. **One idea per sentence.** Short sentences, concrete verbs.
2. **Money with a formula.** "36 breakdowns avoided at a cost difference of €31,000 per breakdown" can be understood and checked; "significant reduction in breakdowns" cannot.
3. **People with a role.** "The doctor reviews and signs the note" instead of "there is human oversight".
4. **Time frames with a date.** "It will be measured by the March session" instead of "soon".
5. **Declared uncertainty.** "We do not yet know whether... We will know when...".
6. **No sales adjectives.** No "revolutionary", "disruptive" or "state-of-the-art".

3.2 Terms best avoided and how to say them

Instead of...	Say...	Why
"The model has an AUC of 0.85"	"Of every 100 customers the system flags as at risk, so many really are; in exchange, it gets so many wrong"	The board needs the effect in terms of hits and errors, not the metric.
"We have launched a PoC"	"We are running a limited three-month trial; in June we will decide whether it continues or stops"	A trial without a decision date is an open-ended cost.
"300% ROI"	"Annual net value of €900,000: efficiencies of €1.1M minus recurring cost of €210,000, validated by management control"	Aggregate ROI hides use cases with negative value and does not say what is validated.
"We have saved 20,000 hours"	"We have released 20,000 hours; of these, 6,000 have been reassigned and none has yet reduced cost"	Released capacity is not a saving until it is materialised.
"Digital transformation with AI"	"This use case does the same thing at lower cost: it is optimisation" or "This use case creates a service that did not exist: it is transformation"	The word is used for different things.

Instead of…	Say…	Why
"Responsible AI"	"Before production we tested whether it treats any group worse; the result is…"	A principle without a verifiable control does not inform.
"Hallucinations"	"Incorrect answers presented as true"	A term that is both technical and colloquial.
"*Prompt* injection"	"Manipulation of the system through instructions hidden in texts or documents"	It describes the risk, not the technique.
"*Human in the loop*"	"A person validates each action before it is executed" (A1) or "The system acts and a person reviews afterwards" (A2)	The difference between the two is what matters to the board.
"Model drift"	"The system loses accuracy over time because reality has changed; it is monitored like this…"	
"Agents" with no further explanation	"A system that, as well as recommending, executes actions: sends communications, modifies data or makes payments"	The risk depends on what it can do.
"Guardrails"	"Limits that prevent the system from doing certain things, for example, granting refunds above a certain amount"	
"LLM", "RAG", "*fine-tuning*"	Avoid them. If essential: "a language model that answers using internal documentation"	They do not change the board's decision.
"We are at the forefront of the sector"	Omit it or cite a verifiable, dated source	Without a source it is an unprovable claim.
"Controlled risk"	"Medium residual risk, with these two controls, tested on such a date"	
"We comply with the AI Act"	"Of 42 inventoried systems, 38 are classified on legal criteria; 3 are high-risk and 2 have their impact assessment"	Compliance is a status per system, not a global statement.

04

Response format for the board

4.1 The four formats

Every board question that admits a yes or no answer, and every decision proposal, is answered with one of these four formats (common specification §5.8):

Format	When it is used	What it must include
"Yes"	The answer is affirmative and there is verified evidence.	The main evidence in one sentence.
"Yes, with one condition: ..."	It is affirmative if something specific is met.	The condition, its owner and its time limit. A single condition; if there are several, the real answer is "not yet".
"Not yet, because ... is missing"	It may become affirmative, but it is not today.	Exactly what is missing, who provides it and when it will be brought back.
"No, because ..."	The answer is negative.	The main reason and, if there is one, the alternative.

Rules:

- **"It depends" is not an answer.** If it depends on something, the format is "Yes, with one condition" or "Not yet, because ... is missing".
- **The condition cannot affect critical controls** for security, legal compliance or human oversight (01 §7.3). If one of those controls is missing, the answer is "Not yet".
- **"Not yet" generates a recorded commitment** with an owner and a date (document 62).
- **"No" is a valid and valuable answer.** Avoiding a bad investment also creates value.

4.2 Examples

All the examples are fictitious.

"Yes"

> **Question:** Can we stop the customer service agent if it starts to fail? **Answer:** Yes. It has a kill switch that diverts all conversations to people; it was tested in the May drill and took four minutes.

> **Question:** Do all the use cases in production have their continuity review up to date? **Answer:** Yes. All 14 have had a review in the last three months, verified by the AI Auditor; the detail is in the portfolio block of the dashboard.

"Yes, with one condition: ..."

> **Question:** Do we extend automated clinical documentation to all consultations? **Answer:** Yes, with one condition: that medical management approves before January the activity to which the released hours will be allocated. Without that reassignment, the extension adds cost and does not add value.

> **Question:** Do we authorise the first stage of the subscription service? **Answer:** Yes, with one condition: that the investment for the stage does not exceed €400,000 and that the decision to continue is brought back in nine months with the contracts signed.

"Not yet, because ... is missing"

> **Question:** Do we put the inspection prioritisation system into production? **Answer:** Not yet, because the fundamental rights impact assessment is missing. Legal counsel will deliver it on 15 November and it will be brought to the December session.

> **Question:** How much value does generative AI deliver in the company? **Answer:** We cannot say yet, because measuring it with a control group is missing. Today we have usage (1,200 weekly users) and cost (€330,000 per year), but not value. Management control will present the method at the next session.

"No, because ..."

> **Question:** Do we let the battery optimisation system operate in the market without oversight? **Answer:** No, because an order executed in the market cannot be reversed and the residual risk is Critical, outside the approved appetite. The alternative is for it to recommend and for an operator to decide.

> **Question:** Can we say that the company is transforming itself with AI? **Answer:** No, because 92% of validated value comes from efficiencies and no Transform use case has reached production. We are efficient, which is a good outcome, but it is not transformation.

4.3 Structure of an intervention

1. **Answer** with one of the four formats (one sentence).
2. **Reason** (one or two sentences).
3. **Evidence**: which document, data or test supports it and who has verified it.
4. **What is asked of the board**: to decide, to take note or nothing.

05

How many use cases to present

5.1 Rule

In each session, **at most three use cases are presented in detail**. Three use cases allow the board to understand, ask questions and decide; with more, the session becomes a presentation and decisions are postponed.

5.2 How to choose them

The selection criteria are those in document 60 (section 4.8). A balanced combination is usually:

Use case	What it gives the board
A use case that requires a decision	The session has a specific purpose.
A use case involving risk, an incident, a stop or a retirement	The board sees that the control system works and learns from what does not go well.
A value or transformation use case	The board follows its bets and the value with evidence.

5.3 How to tell a use case in five minutes

Minute	Content
1	What it is and what it is used for, in business language. Sphere and ambition level.
2	What value it delivers: formula, amount and validation status; recurring cost; net value.
3	What risk it carries: residual level, regulatory classification, who oversees it and how it is stopped.
4	What has happened since the last session and what the relevant decision-maker has decided.
5	What is asked of the board, using the response format.

06

Questions by sphere

The reference questions come from the sphere map (00 §4.1). Each sphere includes follow-up questions, the evidence the board should request and the signal that should draw its attention.

6.1 Spheres where value is created

Sphere	Reference question	Follow-up questions	Evidence to request	Warning signal
01 · Customer	If a competitor knows our customers better because it uses AI better, how long will it take us to lose relevance?	What decisions about customers does a system take or prepare today? Do we know whether it treats any group worse? What part of the customer relationship goes through an assistant?	Use cases by ambition level in the sphere; retention or conversion measured with a control group; fairness tests; rates of referral to people.	All the value declared in the sphere is customer service savings; there is no use case with measured return.

Sphere	Reference question	Follow-up questions	Evidence to request	Warning signal
02 · Product and service	Are we using AI to protect the current product or to build the next one?	Is there any product or service that exists only thanks to AI? Do protecting and building compete for the same resources?	Transform use cases in the sphere with their learning milestones and stage limits; AI-enabled revenue.	No Transform bet approved, or bets approved without stop criteria.
03 · People	Does AI replace, augment or reorganise people's work?	What happens to released capacity? Have we informed employee representatives where required? What new roles do we need?	Position approved by management; released capacity materialised and reassigned; training and AI literacy plan; systems that assess employees and their classification.	Released hours presented as savings; no explicit position on the effect on employment.
04 · Operations	Which processes have a demonstrable short-term return and which require a complete redesign?	Are tasks being automated or are processes being redesigned end to end? What happens if the system goes down?	Baseline and validated value per process; tested rollback plan; incidents and recovery time.	Many pilots without a decision date; no rollback plan in critical processes.

6.2 Enabler spheres

Sphere	Reference question	Follow-up questions	Evidence to request	Warning signal
05 · Data	Do we know what data we have, who is accountable for its quality and whether we can legally use it for AI?	How many use cases have been stopped or delayed because of data? What data do we share with model suppliers and on what terms?	Data owners by domain; stop reasons due to insufficient data; documented legal basis; data use clauses in contracts.	"Insufficient data" is the most frequent stop reason and there is no use case to resolve it.
06 · Knowledge	How much critical knowledge depends on a few people and is not available to the organisation?	If the people who know most left tomorrow, what would be lost? Do internal assistants answer with verifiable sources?	Critical knowledge map; assistants with measurement of correct answers and cited sources.	Internal assistants in use without measurement of the quality of their answers.
07 · Decision	What does AI decide alone, what requires human validation and what is never delegated?	Are there systems that act without individual review? Is there a record of what the system recommended and what the person decided?	Inventory with autonomy level A0–A3; approved list of non-delegable decisions; log of agent actions.	There is no list of non-delegable decisions: decisions are being delegated by default.

6.3 Boundaries and meta-sphere

08 · Regulation, ethics and accountability. Reference question: *Do we govern proactively or do we wait for the regulator to force us?*

Assessment axis	Questions	Evidence to request
Compliance	Are all systems inventoried and classified on legal criteria? Have the required impact assessments been carried out? Do we know how to notify a serious incident on time?	Inventory with regulatory classification; data protection and fundamental rights assessments; notification procedure and drill.
Anticipation	What new obligations will affect us in the next twelve months? Who monitors changes?	Regulatory mapping with consultation and application dates (document 34).
Ethical leadership	Do we have our own principles beyond the legal minimum? Can anyone veto a use case on ethical grounds? Do we inform the public about how we use AI?	Approved corporate policy; body with real veto power and recorded decisions; public information.

Warning signal: the board only hears about regulation when there is a penalty or an incident.

09 · AI governance. Reference question: *Does AI have corporate governance with the same formality as finance, risk or compliance?*

Assessment axis	Questions	Evidence to request
Structure	Who has ultimate responsibility for the AI strategy? Is there a portfolio prioritised by value and risk? Does the board receive regular and understandable information?	Approved AI thesis; AI Committee with a mandate; quarterly pack (document 60).
Balance between speed and control	Are there decision gates with a real ability to stop? How long does it take from idea to approval and to production? Has any use case been stopped?	Register of *gate* decisions; agility metrics by risk level; stop reasons.
Supplier ecosystem	Do we depend on a single model supplier? Do we decide to build, buy or partner on a case-by-case basis? Do the contracts protect our data and our intellectual property? Do we understand the cost structure?	Supplier register with criticality and exit strategy; cost per use case with allocation.

Warning signal: nobody can say when the last use case was stopped, or why.

07

Questions by stage of the corporate cycle

Stage	Board questions	Evidence that answers them
Diagnosis (C1)	How many AI systems do we have, including suppliers' systems and those employees use of their own accord? What is our real maturity and on what evidence? How much value is validated and how much does AI cost today? Are we efficient or are we transforming?	Inventory; D1–D7 maturity with evidence; value by status and recurring cost; transformation index profile.
Direction (C2)	In which spheres do we want to Optimise, Augment or Transform? What risks will we not accept under any circumstances? Above what investment must the board decide? What will we not do with AI?	AI thesis; ambition by sphere; risk appetite and thresholds; corporate policy; framework budget.
Portfolio (C3)	What do we fund and what do we stop? Is the portfolio balanced between efficiency and transformation? Where does the next euro deliver most? Do we have the capacity to execute it?	Prioritised portfolio with additional net value per euro; distribution by ambition level; retirement criteria; available capacity.
Oversight (C4)	Is what was decided being carried out? What has changed since the last session? What incidents have there been and how were they managed? Which recommendations are overdue? What do we have to decide today?	Quarterly pack; dashboard; register of recommendations and decisions.
Review (C5)	Have we progressed in maturity, with evidence? Has our transformation profile changed? Has the thesis been fulfilled? What have we learned from what we stopped and retired? What do we change for the next cycle?	Review of maturity and index; actual versus target ambition; stop and retirement reasons; proposed adjustments.

08

Script for a 60-minute session

8.1 What it is for

This is the script for a **first working session with the board** on AI: it presents the impact map, positions the company and ends with a mandate for the first ninety days. It is used before C1 or as the opening of C2 when the board has not previously worked with the framework.

8.2 Preparation

When	What	Who
Two or three weeks before	Short interviews with the chair, the chair of the board committee and senior business, technology and risk management. Preliminary inventory of AI systems.	The presenter, with the AI Office if there is one
Two weeks before	Selection of one example per ambition level in spheres 01 to 04, ideally from the company itself; if there are none, fictitious sector examples marked as such.	The presenter
One week before	Circulation of a two-page note: session objective, nine spheres, three levels, the reference questions and the decision that will be requested at the end.	Board secretariat
The day before	Timed rehearsal. Each block fits its time or is cut.	The presenter

8.3 Script

Block	Minutes	Objective
01 · Strategic framing	0–10	Position AI as a business issue and present the impact map.
02 · Spheres where value is created	10–22	See where the company's AI is today and where it could be.
03 · Enabler spheres	22–32	Diagnose real capability: without data, knowledge and decision rules, nothing works.
04 · Regulation, ethics and accountability	32–42	Find out whether the company governs proactively or reactively.

Block	Minutes	Objective
05 · Position on people	42–50	Obtain an explicit position from the board on the effect of AI on work.
06 · AI governance and first 90 days	50–60	Close with a decision: mandate, owner and report-back date.

Block 01 · Strategic framing (0–10 minutes)

Opening question	"What business are we in and how does AI change its rules?"
What is presented	Why AI is a board matter: risk, opportunity, accountability and competitive advantage. The nine spheres grouped into where value is created (01–04), enablers (05–07), boundaries (08) and meta-sphere (09). The three ambition levels: Optimise, Augment and Transform, which are not a ladder but three types of bet with different cost, risk, time frame and organisational resistance.
Material	Sphere graphic component; table of ambition levels (00 §4.1).
What the board is asked	Which two or three spheres concern you most today?
Outcome	Priority spheres noted for the rest of the session.
Avoid	Starting with technology or with the evolution of models.

Block 02 · Spheres where value is created (10–22 minutes)

What is presented	Customer, Product and service, People and Operations, with an example of each ambition level per sphere. For example, in Customer: classifying and routing queries (Optimise), personalising the proposal for each customer (Augment), a main relationship channel handled by an assistant with oversight (Transform).
Material	Map of the current portfolio by sphere and level, even if preliminary; if there is no data, say so.
What the board is asked	The reference questions for spheres 01 to 04 (section 6.1). In particular: "Are we using AI to protect the current product or to build the next one?".
Outcome	First reading of where the portfolio stands and whether there is any Transform bet.
Avoid	Going through every possible example. One per level is enough.

Block 03 · Enabler spheres (22–32 minutes)

What is presented	Data, Knowledge and Decision. An honest diagnosis of real capability: quality and legal basis of the data, dependence on the knowledge of a few people, which decisions are already delegated to systems.
Material	Findings from the prior interviews; preliminary inventory with autonomy level if known.
What the board is asked	"Have we defined which decisions AI can take alone, which require human validation and which must never be delegated?"
Outcome	Recognition of the enabling gaps that constrain any ambition.
Avoid	Turning the block into a review of the technology architecture.

Block 04 · Regulation, ethics and accountability (32–42 minutes)

What is presented	What the EU AI Act requires according to the risk level of each system; data protection in automated decisions; applicable sector regulation. The three assessment axes: Compliance, Anticipation and Ethical leadership. It is often the topic that concerns the board most.
Material	Summary regulatory mapping with consultation date; status of the inventory and classification.
What the board is asked	"Do we govern proactively or do we wait for the regulator to force us?"
Outcome	Knowing whether there is an inventory with regulatory classification and who is accountable for it.
Avoid	Presenting regulation as a catalogue of penalties or treating as settled interpretations that require legal judgement.

Block 05 · Position on people (42–50 minutes)

Question	"Will AI replace, augment or reorganise people's work in this company?"
What is presented	The three simultaneous roles of employees with regard to AI: **recipients** (their tools and expectations change, and adoption depends on change management), **subjects** (AI may restructure or eliminate their role, and the board is accountable for how this is managed) and **promoters** (whoever knows a process best identifies uses that no outsider would see).
What is asked of the board	An explicit position, even if provisional. If the board is not in a position to set it, a proposal for C2 is commissioned with a date.

Outcome	Position or assignment recorded.
Avoid	Watering down the question to make it comfortable. That answer defines the culture, the ability to attract and retain talent and the relationship with the workforce for years.

Block 06 · AI governance and first 90 days (50–60 minutes)

What is presented	The meta-sphere: Structure, Balance between speed and control and Supplier ecosystem. The ninety-day plan in 01 §5.3: **month 1**, maturity diagnosis with evidence and inventory; **month 2**, map of risks and opportunities by sphere with owner, economic impact and time frame; **month 3**, governance structure with bodies, roles, thresholds, *gates*, metrics and reporting rhythm.
What is asked of the board	A decision using the format in section 4: "Do we commission the diagnosis and the governance structure within ninety days?".
Outcome	Recorded decision (DEC) with an executive owner, the body that receives the result and the date of the report-back session.
Avoid	Ending without a decision or with "we will look into it".

8.4 After the session

Time limit	Action	Accountable
2 working days	Note of conclusions: priority spheres, position on people, decision adopted.	The presenter
5 working days	Recording of decisions and assignments with code, owner and date (document 62).	Board secretariat
Next session	First item on the agenda: status of what was decided.	Chair of the AI Committee

8.5 Variant for oversight sessions (C4)

Minutes	Content	Supporting document
0–5	Main message using the response format; changes since the previous session.	One-page summary (60 §4.2)
5–20	Decisions requested, one by one.	Decision sheets (60 §5)

Minutes	Content	Supporting document
20–35	Three use cases in detail, five minutes each.	Section 5.3 of this guide
35–45	Risks, incidents and nonconformities; programme traffic lights.	60 §4.5, §4.6 and §9
45–55	Overdue recommendations and discrepancies between declared status and assessment.	Register (document 62)
55–60	Summary of decisions and assignments, with owner and date.	—

09

Frequent difficult questions

The answers are indicative and are adapted using the company's data. All figures in the examples are fictitious.

#	Question	How to answer	Supporting evidence	What to avoid
1	**How much does AI save us?**	Separate efficiencies, return and recurring cost; give the net value and the validated proportion; state unmaterialised released capacity separately.	Value tracking per use case with status (T12).	A single figure without status or one that adds up released hours.
2	**Are we behind the competition?**	"We do not have a reliable comparison" if there is no verifiable source. Offer the comparison that is reliable: the company against itself over time (maturity, index, agility).	Maturity and transformation index with evidence; external sources only if verifiable and dated.	Citing market figures without a source.
3	**Why don't we use agents like others do?**	Answer by sphere and risk: where it would make sense for a system to act, with what level of autonomy and what controls would be needed. If there is a proposal, take it through the cycle.	Inventory with autonomy A0–A3; agent security criteria (document 35).	Adopting a technology so as not to appear to be lagging behind.

#	Question	How to answer	Supporting evidence	What to avoid
4	**How many jobs will AI eliminate?**	Give the approved position (replace, augment or reorganise) and the data: released, materialised and reassigned capacity; roles affected; training plan. If there is no position, say so and propose setting it in C2.	Adoption and capacity plan (T20); document 50.	Denying the effect or giving unfounded figures.
5	**Do we comply with the AI Act?**	Answer by system: how many are inventoried, how many classified on legal criteria, how many are high-risk and which assessments have been completed or are pending. State the date on which the regulation was consulted.	Inventory and classification (T02, T07); regulatory mapping (document 34).	"Yes, we comply" without a complete inventory.
6	**Who is accountable if AI gets it wrong?**	Name the roles: the AI Sponsor is accountable for the value and the investment; the AI Operations Owner, for stability and incidents; the AI Risk Owner, for assessment and monitoring; the company is accountable to third parties. Accountability is not transferred to the system or the supplier.	Role assignment record (P03); contracts with suppliers.	"It is the supplier's responsibility".
7	**Why have we stopped this after investing?**	"Because the stop criterion set before investing was not met". Give the investment consumed, the investment avoided and what was learned.	*Gate* decision record with stop criteria (P29).	Making excuses or presenting the stop as a failure of the team.
8	**Why does it take us so long to put anything into production?**	Give the times per stretch (idea → approval → production) and by risk level; point out the bottleneck and what is proposed.	Funnel metrics (03 §3.5).	Proposing to skip *gates* without analysing where the delay lies.

#	Question	How to answer	Supporting evidence	What to avoid
9	**What happens if the system fails or we are attacked?**	Explain the severity of possible incidents, containment, the kill switch or the tested rollback plan, and the notification time limits.	Incident response plan (P26); rollback test; drills.	Claiming that it cannot happen.
Sphere map and ambition levels (10)	**Do we depend too much on one supplier?**	Give the concentration: which use cases depend on each supplier, their criticality and whether there is a tested exit strategy or only a written one.	Supplier register (T09); requirement level N1–N3.	"It is the market leader" as an answer.
Maturity model (11)	**Is this transformation or just efficiency?**	Apply the five classification questions (00 §5.2) in the room. If the first four are negative, it is optimisation, and it is fine for it to be so.	Confirmed ambition classification (P07).	Calling an automation transformation.
Transformation index (12)	**Why is validated value so low?**	Explain what is missing for validation (baseline, control group, management control validation) and the plan with dates. A low, honest percentage is preferable to a high, unsupported one.	Value tracking by status (T12).	Reclassifying declared amounts as validated.
AI thesis, ambition and risk appetite (13)	**How much does AI really cost?**	Give the recurring cost by category and by use case, including model consumption and operations people. If there are uncharged costs, say so.	Cost per use case with allocation (T13; document 42).	Presenting only the licences.
Portfolio management (14)	**Are employees using unauthorised AI tools?**	Give what is known: detections, active technical controls, training, regularisations. If it is not measured, "Not yet, because … is missing".	Corporate use monitor (T21); nonconformities for unauthorised use.	Saying it does not happen because it is prohibited.
15	**Can we speed up if we relax the controls?**	"Yes, with one condition" only for non-critical controls and with acceptance of the risk by the appropriate person; "No" for security, legal compliance and human oversight.	Risk appetite approved in C2; rules on conditions (01 §7.3).	Giving in to deadline pressure without a record.

#	Question	How to answer	Supporting evidence	What to avoid
16	**What is stopping us from scaling AI?**	Name the three barriers to start with, with what they block, what will be done, who and the annual value at stake; separate what blocks a gate from what only holds back value.	«What holds back scaling» reading of the dashboard (document 60 §10.4) and the decisions or recommendations that turn it into actions (document 62).	A list of problems without order or owner, or presenting the value at stake as secured value.
17	**How do we know it is worth it if it cannot be put in euros?**	Separate the three perspectives: what does yield in euros (with its status), what changes in operations (indicator with baseline, target and current value) and the risk taken on; if the use case is sustained by non-quantified value, say when it will be reviewed again.	Presentation rule 11 and section 5.3 of document 40; use case record in the dashboard and its next R6.	Converting the improvement into euros without a formula ("worth a million in brand value"), or defending the use case only with the team's perception.

10

Role of the director or adviser with AI experience

10.1 What it is for

SEVEN-G provides for the board or its board committee to be able to rely on a **director or adviser with AI experience** (01 §8.3). Their function is to help the board **ask the right questions, interpret the evidence and follow up on what has been decided**, not to manage the company's AI.

10.2 What they do and do not do

Does	Does not
Helps prepare the agenda and the board's questions.	Lead the AI Office or the AI Committee.
Reviews the pack before the session and points out figures without status, undeclared omissions and technical jargon.	Prepare the pack that they then review.
Translates into business language the technical issues raised by management or the board.	Take decisions that belong to the board, the committee or the sponsors.

Does	Does not
Makes recommendations to the board, which are recorded with a code (document 62).	Assign themselves the assessment of recommendations whose execution depends on work that they or their organisation provide.
Provides the independent assessment of compliance with recommendations when the board commissions it.	Assess as met anything that lacks evidence.
Contributes to the board's training in AI.	Recommend suppliers, products or services in which they have an interest.

10.3 Independence

- **Declaration of interests** at the start of the mandate and whenever it changes: relationships with technology suppliers, consultancies and the company itself (P41, Model C).
- **Incompatibility** with providing implementation services on the same use cases they oversee (document 91, section 5).
- **Access to information** through the board secretariat and formal channels, not through direct relationships with teams that bypass management.
- **Confidentiality**: the company's information is not used in any other context.

10.4 Board training

The EU AI Act includes AI literacy obligations for the staff of providers and deployers (Article 4). Regardless of their exact scope for governing bodies, which must be assessed on legal criteria, the board should receive sufficient training to exercise its oversight. An initial session using the script in section 8 and an annual update with the C5 review are a reasonable basis. Regulatory references consulted in September 2026; their currency must be verified. *This document does not constitute legal advice.*

11

Frequent mistakes

Mistake	Consequence	What to do
Presenting an inventory of technical projects	The board can neither prioritise nor hold to account.	Present spheres, ambition levels, value and risk.
Bringing ten use cases to the session	Nothing is decided and the session drags on.	Three use cases in detail; the rest in the dashboard.

Mistake	Consequence	What to do
Giving figures without a validation status	Decisions about value that does not exist.	Status on every amount and validated proportion visible.
Adding released hours as savings	Inflated value that does not appear in the income statement.	Report released capacity separately.
Answering "it depends"	The board does not know what to decide.	Use one of the four formats.
Hiding stops and incidents	Loss of trust when they become known.	Present them with what was learned.
Calling any use case transformation	The company believes it is transforming and does not take the decisions that doing so would require.	Classify with the five questions.
Ending the session without recorded decisions	What was discussed depends on the memory of those present.	Record decisions and assignments with code, owner and date.
Having whoever presents a use case assess its compliance	Self-assessment without credibility.	Independent assessment in the register.
Comparing with the market without a source	Claims that do not withstand a single question.	Compare with oneself over time or cite verifiable sources.
Talking about models instead of decisions	The board disengages.	Translate using the table in section 3.2.
Requesting approval without options	The board cannot perform its function.	Include at least two options, one of them not doing it or stopping.

12

Associated tools and templates

Name	Use in this document
Ambition classifier `(T05)`	Difficult question 11; classification in the room.
Maturity diagnosis `(T15)`	Block 06 of the script; stage C1.
Portfolio sphere map `(T16)`	Blocks 02 and 03 of the script.
Board AI dashboard `(T17)`	Use cases in detail and oversight sessions.

Name	Use in this document
Board recommendations register (T18)	Recording of decisions and assignments after each session.
Sphere and ambition classification (P07)	Evidence of the classification.
***Gate* decision record** (P29)	Answers about stops and criteria.
Use case sheet (P31)	How to tell a use case in five minutes.
Declarations of independence and conflicts of interest (P41)	Declaration of interests by the director or adviser with AI experience (Model C, section 10.3).
Pre-session note and note of conclusions for the board session (P68)	Pre-session note, agenda and conclusions of the session (sections 8.2 to 8.4).
Board decisions register (P69)	Recording of decisions and assignments after the session (section 8.4).

13

Related documents

Document	Relationship
00 · What SEVEN-G is and how it helps companies	Spheres, ambition levels, reference questions and measurement rules.
01 · Foundational methodology	Corporate cycle, bodies and decision rules.
10 · Sphere map and ambition levels	Development of the spheres and their indicators.
13 · AI thesis, ambition and risk appetite	Board decisions in C2.
34 · Regulatory mapping	Basis for regulatory answers.
35 · AI and agent security	Autonomy, agent controls and exposure to attacks.
50 · People and organisation	Position on people and released capacity.
60 · Board pack	Information delivered before the session.
62 · Recommendations and decisions register	Record of what is decided and assigned.
90 · Implementation guide	First ninety days.

Document	Relationship
91 · Guide for consultants	Independence of the adviser.

14

Version control

Version	Date	Changes
0.1	16-09-2026	First version. Integrates the questions and session script from the sphere map into the SEVEN-G framework; defines business language, the response format for the board with examples, the three-use-case rule, the questions by sphere and by stage, the difficult questions, the role of the director or adviser with AI experience and the frequent mistakes.
0.2	28-09-2026	Difficult question 16 «What is stopping us from scaling AI?», answered with the dashboard's «What holds back scaling» reading (60 §10.4).
0.3	01-10-2026	Difficult question 17 «How do we know it is worth it if it cannot be put in euros?», answered with the three perspectives of document 40 §11.1, rule 11 (D150).

External sources cited

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

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