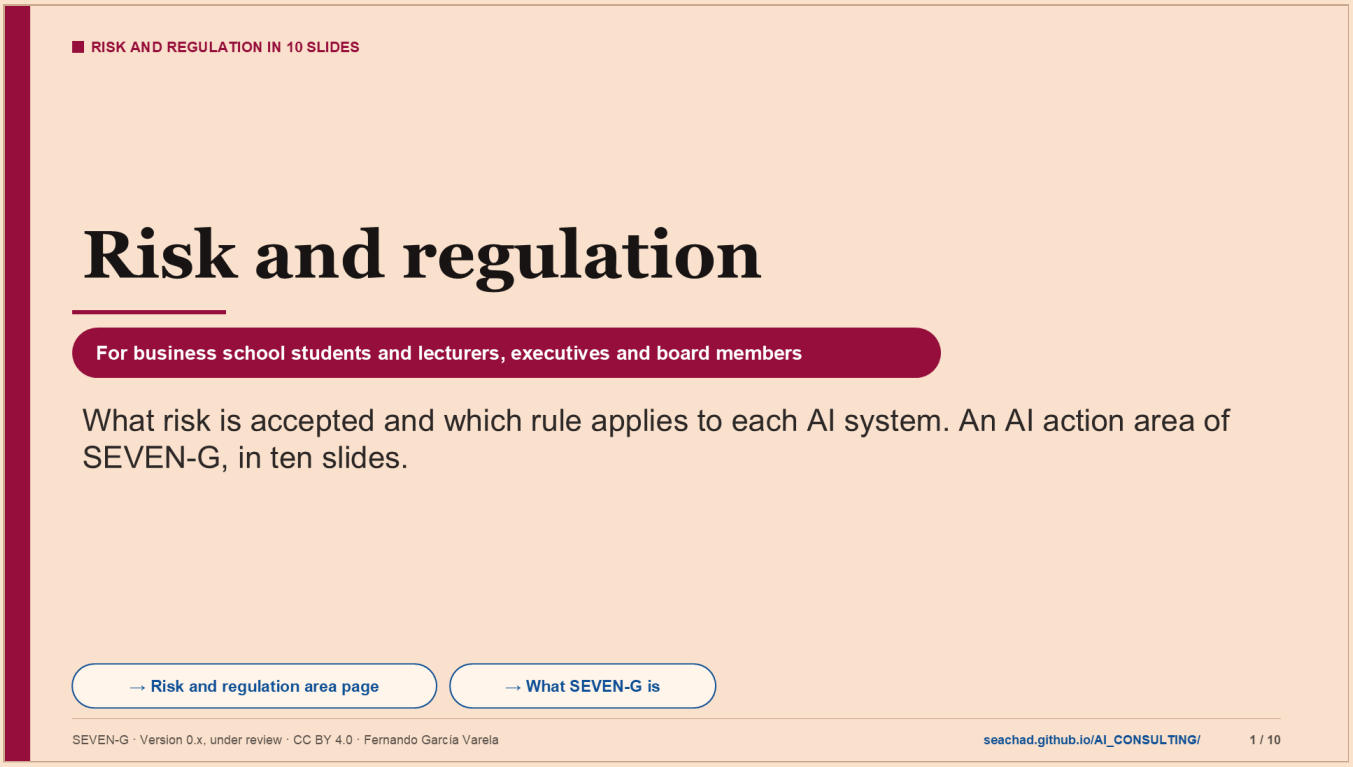

Notes

This presentation explains the Risk and regulation action area of SEVEN-G. SEVEN-G is an open and free framework for a company to govern, measure and decide on AI. Its author is Fernando García Varela and it is published under a CC BY 4.0 licence.

The area answers a double question: what risk the company accepts with each AI system and which rule applies to it. To answer it, SEVEN-G asks for four things:

- Inventory all AI systems, including third-party systems and those used without authorisation.
- Classify each system under the regulation that applies to it, starting with the EU AI Act.
- Assess its risks before and after controls, and have the body that matches its level accept the remaining risk.
- Record incidents and nonconformities and correct them within a deadline.

The ten slides follow this order: why it matters, the key concepts, how SEVEN-G handles it across the lifecycle, the tools and templates, an example with fictitious data, common mistakes, questions for management or the classroom, the relationship with other areas and with SPHERES, and how to start.

A warning that applies to the whole presentation: SEVEN-G is not legal advice and does not guarantee compliance with any rule. Each organisation identifies the regulation that applies to it, including that of its sector, and verifies that it is in force with qualified advice.

Question to open the session: if a supervisor asked you tomorrow for the list of your company's AI systems with their classification, how long would it take you to provide it?

SOURCES AND LINKS

- **Risk and regulation area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/riesgo.html
- **Document 34 §1 · Scope and responsibility of the regulatory mapping** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html#purpose-and-scope
- **What SEVEN-G is** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/entrada/index.html

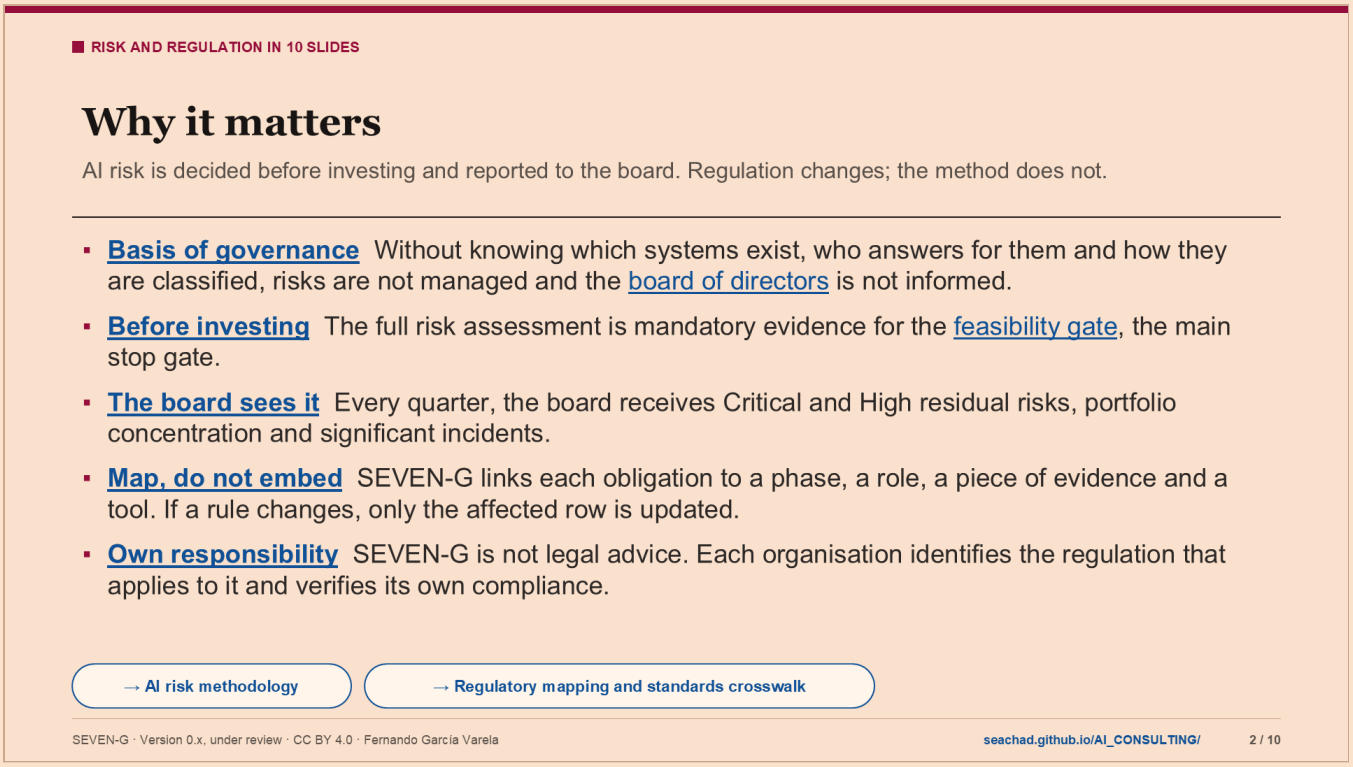

Notes

This slide gives five reasons to take the area seriously.

- Basis of governance. Document 32 calls the inventory the basis of all governance. If the company does not know which AI systems it uses, it does not know which risks it has or which obligations it must meet either.
- Before investing. The first principle of the risk methodology (document 33) is that risk is assessed before investing. The full assessment is required at the feasibility gate (G3). It is the main stop gate: it decides whether to continue before the large spend on design and build.
- The board sees it. Every quarter, the board or its board committee receives Critical and High residuals by sphere, portfolio concentration and scenarios, compliance with the risk appetite and significant incidents. The report uses business language: what can happen, whom it affects, what is being done and what decision is requested.
- Map, do not embed. SEVEN-G does not copy regulation into its phases. It links each obligation to a phase or stage, a responsible role, a piece of evidence and a tool. Why it matters: AI regulation changes quickly, and in this way a regulatory change does not force a rebuild of the lifecycle, the gates or the templates.
- Own responsibility. The regulatory mapping in document 34 is indicative and reflects the state of the rules at a consultation date. The user organisation is solely responsible for identifying the regulation that applies to it, classifying each system and verifying its compliance with qualified legal advice.

Question for the classroom: what costs a company more, stopping an AI use case before building it or discovering after go-live that it was wrongly classified?

TERMS ON THIS SLIDE

- **Board of directors**: The company's highest body. In SEVEN-G it approves the AI thesis and risk appetite (C2), Transform bets and their scaling, and oversees AI every quarter. AI is a strategic asset: its big decisions are not only

for technology. The link opens its role in the governance model (document 30). https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/30_SEVEN-G_Modelo_de_gobierno.html#3-1-board-of-directors

- **Feasibility gate** (G3): The most important gate: using the feasibility test, the regulatory classification and the risk register, it decides whether the use case deserves the big spend of building it. It is always decided on its own, never grouped with others. It is the main stop gate: stopping here costs little; later, a lot. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-4-g3-feasibility

SOURCES AND LINKS

- **Document 33 §2 · Principles of the risk methodology** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#principles
- **Document 33 §12 · Risk reporting** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#risk-reporting
- **Document 34 §1.1 · SEVEN-G maps, it does not embed** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html#1-1-principle-seven-g-maps-it-does-not-embed
- **AI risk methodology** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html
- **Regulatory mapping and standards crosswalk** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html

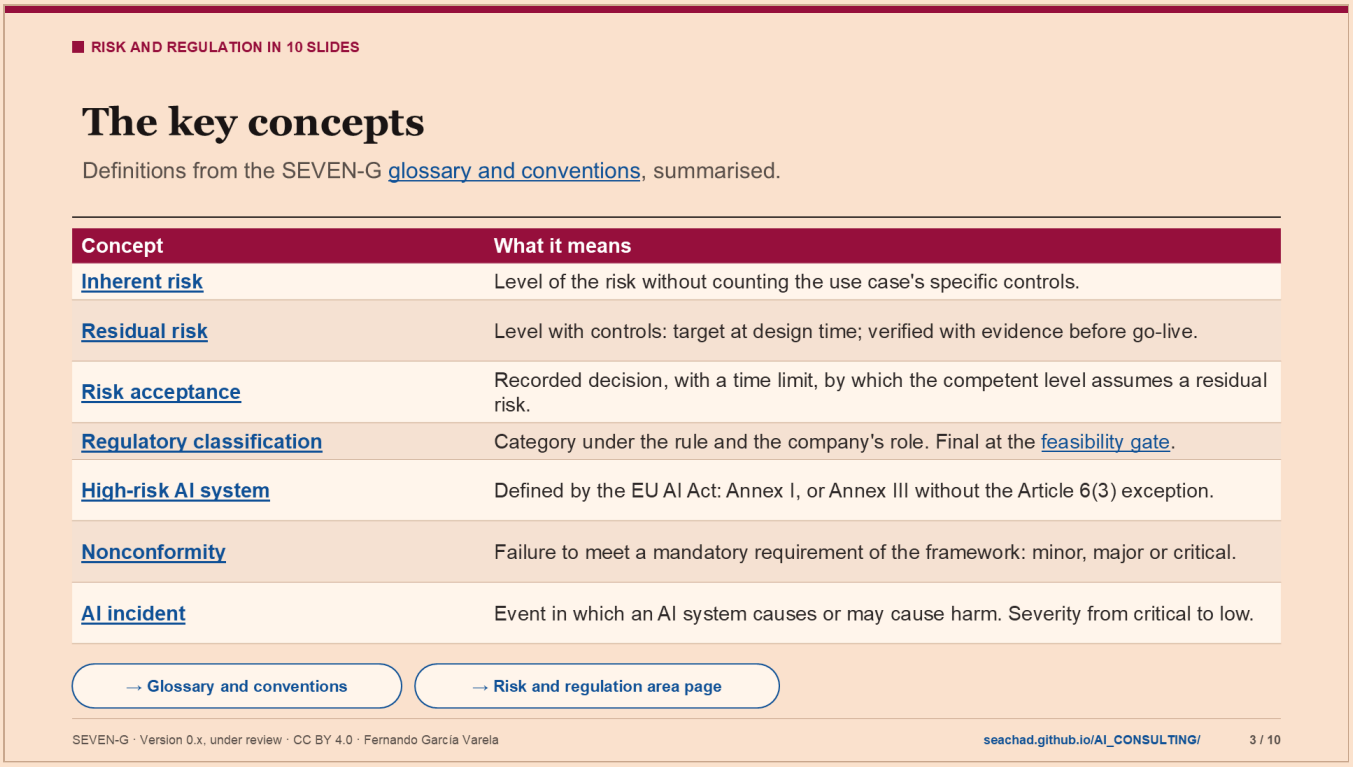

RISK AND REGULATION IN 10 SLIDES

The key concepts

Definitions from the SEVEN-G glossary and conventions, summarised.

Concept	What it means
Inherent risk	Level of the risk without counting the use case's specific controls.
Residual risk	Level with controls: target at design time; verified with evidence before go-live.
Risk acceptance	Recorded decision, with a time limit, by which the competent level assumes a residual risk.
Regulatory classification	Category under the rule and the company's role. Final at the feasibility gate.
High-risk AI system	Defined by the EU AI Act: Annex I, or Annex III without the Article 6(3) exception.
Nonconformity	Failure to meet a mandatory requirement of the framework: minor, major or critical.
AI incident	Event in which an AI system causes or may cause harm. Severity from critical to low.

→ Glossary and conventions → Risk and regulation area page

SEVEN-G · Version 0.x, under review · CC BY 4.0 · Fernando García Varela seachad.github.io/AI_CONSULTING/ 3 / 10

Notes

These seven concepts are the minimum vocabulary of the area. All of them come from the SEVEN-G glossary (document 02).

- Inherent risk. It is the risk before the use case's own controls. General company controls only count if they are tested and documented.
- Residual risk. It is the risk with controls. It has two moments: target, expected after applying the planned controls and used at the feasibility (G3) and design (G4) gates; and verified, when there is evidence that the controls work, required at go-live (G5) and monitored at the continuity review (R6). Why it matters: it separates what is promised from what has been checked.
- Risk acceptance. It is a recorded decision, with a validity period and conditions. Acceptance lapses if the level, system, supplier, model, regulatory classification or autonomy changes.
- Regulatory classification. It determines, with qualified legal judgement, the system's category (prohibited, high risk, transparency obligations, minimal risk, out of scope or pending classification) and the company's role: for example, provider or deployer.
- High-risk AI system. The EU AI Act defines it, not SEVEN-G. It is a criterion that puts the use case into Enterprise intensity, the most demanding one. The application dates of its obligations are dealt with in document 34.
- Nonconformity. It is a failure to meet a mandatory requirement of the framework. It is managed with containment, root cause, corrective action, effectiveness verification, closure and re-audit.
- AI incident. It is an event with actual or possible harm, even if the system remains available. It is classified by severity: S1 critical, S2 high, S3 medium and S4 low.

The EU AI Act, Regulation (EU) 2024/1689 as amended, is the main reference for classification. Its official text is on EUR-Lex.

Question for the classroom: why do you think SEVEN-G separates an incident, which is an event with harm, from a nonconformity, which is a failure to meet a rule?

TERMS ON THIS SLIDE

- **Glossary and conventions** (documento 02): The official SEVEN-G glossary: each term with its definition, its Spanish equivalent and the document that develops it, plus the scales and codes used across the framework. It ensures everyone calls the same thing by the same name. The link opens the document. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/02_SEVEN-G_Glosario_y_convenciones.html
- **Feasibility gate** (G3): The most important gate: using the feasibility test, the regulatory classification and the risk register, it decides whether the use case deserves the big spend of building it. It is always decided on its own, never grouped with others. It is the main stop gate: stopping here costs little; later, a lot. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-4-g3-feasibility

SOURCES AND LINKS

- **Document 33 §5.1 · Inherent and residual risk** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#5-1-definitions
- **Document 37 §1.1 · Definitions of incident and nonconformity** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#1-1-definitions
- **Regulation (EU) 2024/1689 (EU AI Act), EUR-Lex** https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
- **Risk and regulation area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/riesgo.html

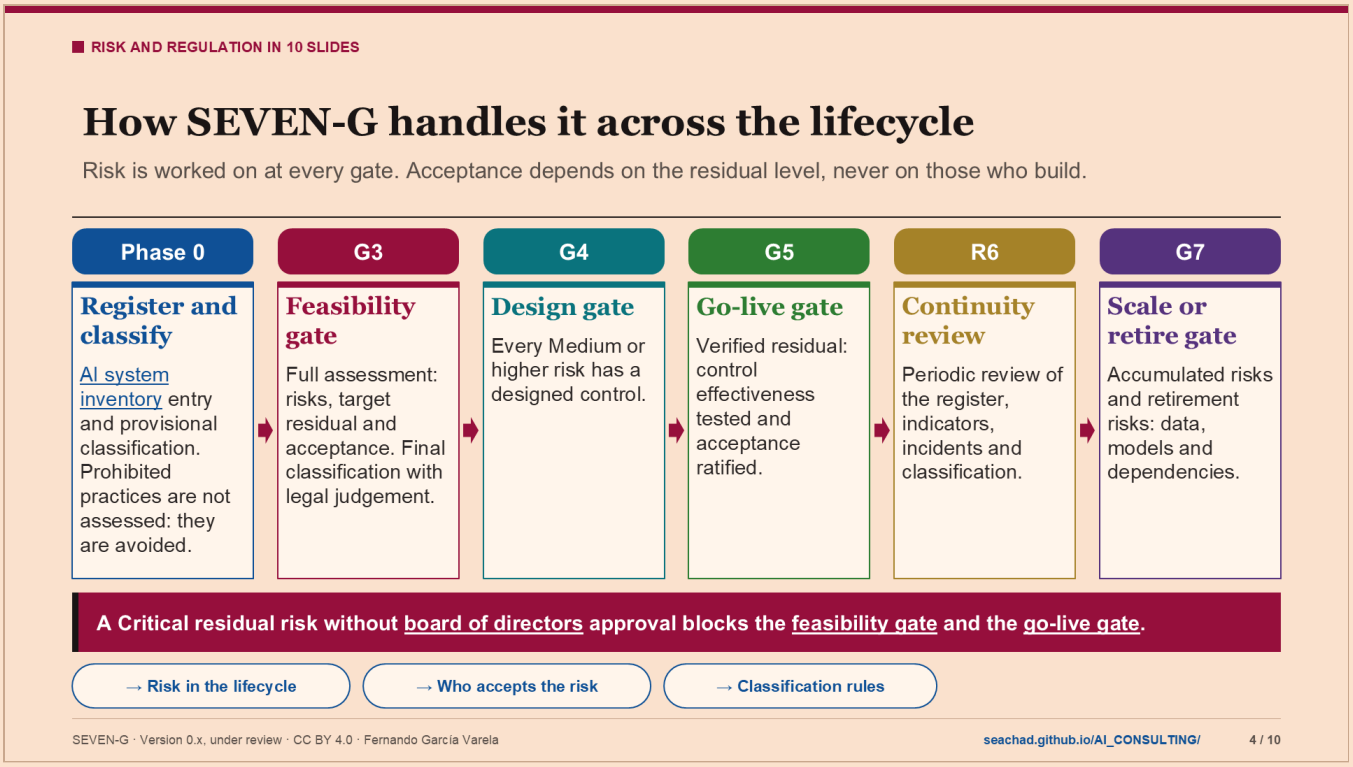

Notes

SEVEN-G does not treat risk as a one-off formality. It works on it at every moment of the use case lifecycle (phases 0 to 7) and of the corporate cycle (C1 to C5).

- Phase 0. The system is registered in the inventory and receives a provisional classification. Obvious risks are identified: prohibited practice, protected data or critical function. A practice prohibited by the EU AI Act is not assessed: it is avoided.
- Feasibility gate (G3). This is the full assessment: inherent risk, planned controls, target residual and acceptance. The classification becomes final, with qualified legal judgement.
- Design gate (G4). Every Medium or higher risk has a designed control.
- Go-live gate (G5). Control effectiveness is tested and the residual becomes verified. Acceptance is ratified.
- Continuity review (R6). The register, key risk indicators, incidents, changes and classification are reviewed.
- Scale or retire gate (G7). Accumulated risks and retirement risks are examined.

Who decides. Acceptance of residual risk follows a fixed scale: Low, the AI Product Owner; Medium, the AI Sponsor with clearance from the AI Risk Owner; High, the AI Committee, with a treatment plan and validity until the next continuity review; Critical, not accepted: only, by exception, the board or its board committee approves it, within the approved risk appetite. No one accepts a risk of a use case they are building.

One rule worth underlining: the extreme impact rule. A risk with impact 5 on the people and rights axis or on the regulatory axis is treated as at least High for acceptance purposes, even if its likelihood is 1. In this way, a rare but unacceptable event is not accepted by the product owner.

Question for the classroom: why do you think SEVEN-G requires the final classification at G3 and not at the end, at G5?

TERMS ON THIS SLIDE

- **AI system inventory** (T02): A list of all the company's AI systems: its own, third-party AI embedded in processes, the assistants employees use and any unauthorised use detected. Each system has an owner, a regulatory classification, an intensity and an autonomy level. It comes first, because you cannot govern what you do not know. It is a view of the use case register: open the link and choose «Inventory» (fictitious data). https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html
- **Board of directors**: The company's highest body. In SEVEN-G it approves the AI thesis and risk appetite (C2), Transform bets and their scaling, and oversees AI every quarter. AI is a strategic asset: its big decisions are not only for technology. The link opens its role in the governance model (document 30). https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/30_SEVEN-G_Modelo_de_gobierno.html#3-1-board-of-directors
- **Feasibility gate** (G3): The most important gate: using the feasibility test, the regulatory classification and the risk register, it decides whether the use case deserves the big spend of building it. It is always decided on its own, never grouped with others. It is the main stop gate: stopping here costs little; later, a lot. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-4-g3-feasibility
- **Go-live gate** (G5): The gate that authorises using the AI system with real users, customers or data. It requires validation and pilot passed, a tested rollback plan and multi-level sign-off (business, technology, risk). Nothing goes into production without this sign-off. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-6-g5-go-live

SOURCES AND LINKS

- **Document 33 §3.1 · Risk in the use case lifecycle** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#3-1-in-the-use-case-lifecycle
- **Document 33 §4.3 · Risk level and 5 × 5 matrix** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#4-3-risk-level-and-5-5-matrix
- **Document 33 §7.1 · Who accepts residual risk** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#7-1-who-accepts-residual-risk
- **Document 32 §5.3 · Classification application rules** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#5-3-application-rules
- **Risk in the lifecycle** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#3-1-in-the-use-case-lifecycle
- **Who accepts the risk** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#7-1-who-accepts-residual-risk
- **Classification rules** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#5-3-application-rules

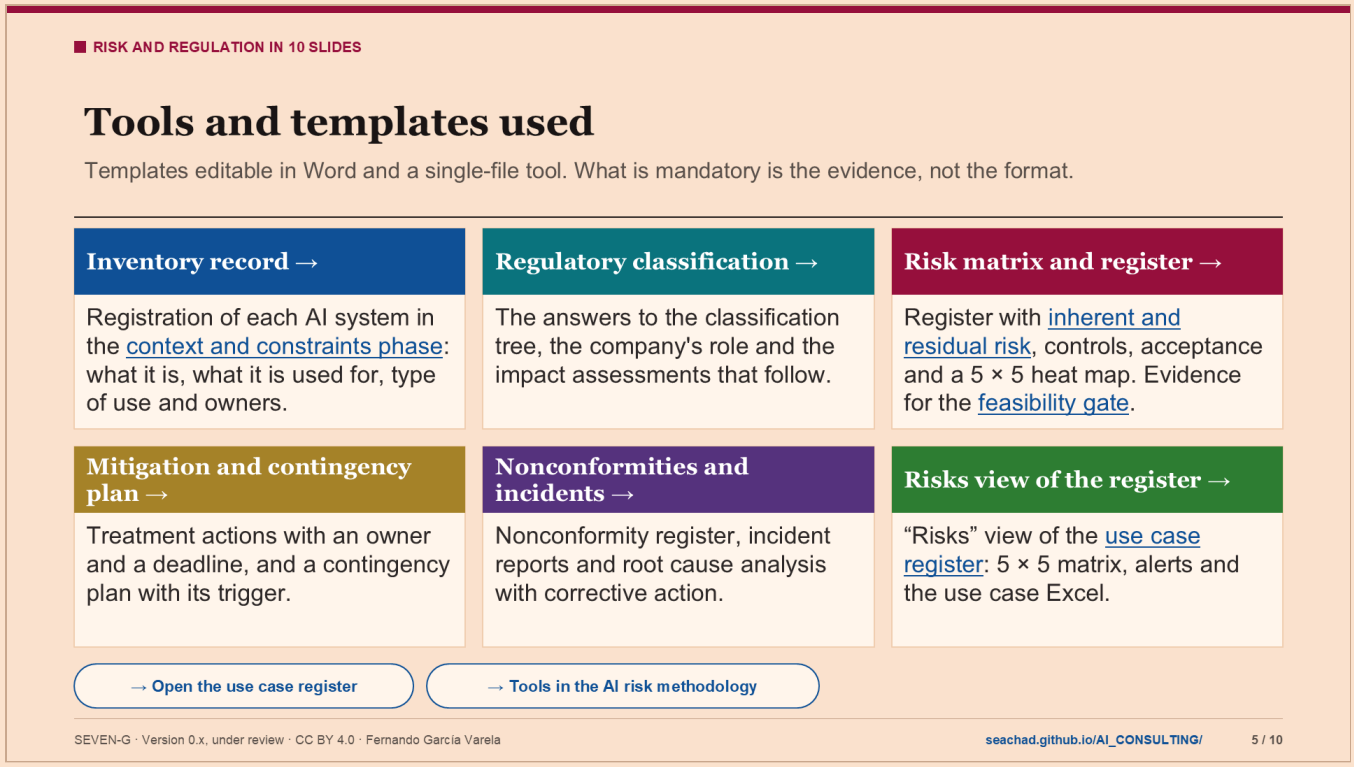

Notes

These are the working pieces of the area. All of them are open and can be adapted to the company's own methodology.

- AI system inventory record (P05). It is the evidence of registration in phase 0. It records an understandable explanation of the system, its type of use and its named owners.
- Regulatory classification and impact assessments (P11). It records, step by step, the answers to the classification tree in document 32, the company's role and which assessments follow. When required, it is completed with the data protection impact assessment (P47) and the fundamental rights impact assessment (P48).
- Risk matrix and register (P12). It is the register with the fields of document 33: typical risk from the catalogue, likelihood, impact by axis, controls and their effectiveness, target and verified residual, response, acceptance and monitoring. It is mandatory evidence for the feasibility gate (G3).
- Mitigation and contingency plan (P13). Every High or Critical risk under treatment must have a contingency plan: what is done if it materialises, who activates it and with what trigger.
- Nonconformities and incidents (P50, P51 and P52). Nonconformity register, incident decisions and reports, and root cause analysis worksheet.
- Risk matrix and register (T06). It is the "Risks" view of the use case register (T01). It calculates levels, draws the inherent and residual 5 × 5 matrix, warns of expired or insufficient acceptances and generates the use case Excel with the matrix and proposed mitigation and contingency plans. The proposals are a starting point: they do not replace the decision of the competent body.

Why it matters: with a common template, all use cases are assessed on the same scale. That way, the AI Committee and the board can compare and add up. SEVEN-G also accepts evidence the company already has, if it covers what the template asks for.

Question for the classroom: what is the advantage of using the same risk scale for AI as the rest of the company's corporate risk framework?

TERMS ON THIS SLIDE

- **Context and constraints** (fase 0): The enabling phase of the lifecycle. Before talking about value, the use case sets its context: sponsor, roles, governance intensity, inventory entry and limits (data, rules, budget). The authorisation gate (G0) closes it. It avoids starting a use case without an owner or rules. The link opens the phase manual in document 20. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/20_SEVEN-G_Manuales_de_fase.html#phase-0-context-and-constraints
- **Inherent and residual risk**: Inherent risk is the risk a use case has without controls; residual risk is what remains after applying them. Both are measured as likelihood by impact on a 5 × 5 matrix. The residual level decides who must accept the risk: only the board accepts a Critical one. The link opens the explanation in document 33. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#inherent-risk-controls-and-residual-risk
- **Feasibility gate** (G3): The most important gate: using the feasibility test, the regulatory classification and the risk register, it decides whether the use case deserves the big spend of building it. It is always decided on its own, never grouped with others. It is the main stop gate: stopping here costs little; later, a lot. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-4-g3-feasibility
- **Use case register** (T01): It is the central SEVEN-G tool and the single data source for all the others. It records each AI use case like an opportunity in a CRM: phase, status, dates, criteria for each gate, risks, value and stop reason. The AI office and the teams use it; the committee and the board see what comes out of it. The link opens an example register with fictitious use cases; anything you change is stored only in your browser. https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html

SOURCES AND LINKS

- **Document 33 §13 · Associated tools and templates** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#associated-tools-and-templates
- **Document 32 §10 · Associated tools and templates** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#associated-tools-and-templates
- **Template P12 · Risk matrix and register** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P12_SEVEN-G_Matriz_y_registro_de_riesgos.html
- **Document 37 §10 · Associated tools and templates** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#associated-tools-and-templates
- **Tools in the AI risk methodology** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#associated-tools-and-templates

RISK AND REGULATION IN 10 SLIDES

An example with fictitious data

Customer service conversational assistant (IA-2025-002), in production, from the example use case register. Fictitious data.

Risk	Without controls	With tested controls	Who accepts it
A prompt injection reveals another customer's order data	4 × 4 = 16 · Critical	2 × 4 = 8 · Medium	Use case sponsor, with clearance from the AI risk owner
Incorrect answers on the returns policy	3 × 4 = 12 · High	2 × 4 = 8 · Medium	AI Sponsor, with clearance from the AI Risk Owner
Customers do not know they are talking to an AI	2 × 3 = 6 · Medium	1 × 3 = 3 · Low	AI Product Owner

Classification: transparency obligations. Enterprise application intensity because of direct exposure to customers.

→ See the example in the register | → Risk assessment scales

SEVEN-G · Version 0.x, under review · CC BY 4.0 · Fernando García Varela seachad.github.io/AI_CONSULTING/ 6 / 10

Notes

This example comes from the site's example use case register (T01). The company and all the data are fictitious.

The system is a generative AI assistant on the website and the app that answers customer questions about orders and returns and hands over to a person when it cannot resolve them. It is in production (phase 6).

Its regulatory classification is “transparency obligations”: people talk to an AI system and must know it. It is not high risk. Even so, it has Enterprise intensity, because it has direct exposure to customers, which is one of the Enterprise criteria. The data protection impact assessment is done; the fundamental rights impact assessment does not apply.

How to read the three rows:

- Prompt injection. A customer writes a message that manipulates the assistant into revealing another customer's data. The determining axis is people and rights. Without controls it is likely (4) and major (4): 16, Critical. The controls are session isolation per customer, an output filter for personal data and adversarial tests before G5 with no leaks. As they are effective and tested, they lower the likelihood by two levels: 2 × 4 = 8, Medium.
- Incorrect answers. The determining axis is reputational. The controls are answers grounded in a document base and handover to a person. It goes from High to Medium.
- Transparency. The determining axis is regulatory. The control is a notice at the start of each conversation and in the interface. It goes from Medium to Low.

Note two things. First, impact does not fall in the first two rows: the controls mainly reduce likelihood. Second, who accepts depends on the residual level: Medium is accepted by the AI Sponsor with clearance from the AI Risk Owner; Low, by the AI Product Owner.

Question for the classroom: if tomorrow an adversarial test achieved a leak, what would happen to the control's effectiveness, to the verified residual and to the acceptance?

TERMS ON THIS SLIDE

- **Use case register** (T01): It is the central SEVEN-G tool and the single data source for all the others. It records each AI use case like an opportunity in a CRM: phase, status, dates, criteria for each gate, risks, value and stop reason. The AI office and the teams use it; the committee and the board see what comes out of it. The link opens an example register with fictitious use cases; anything you change is stored only in your browser. https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html
- **Use case sponsor**: The business person who asks for the use case, answers for its value and defends its budget. Without a sponsor, the use case does not pass the authorisation gate. They cannot verify or audit their own use case. It matters because AI without a business owner stays a pilot. The link opens the roles in document 01. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/01_SEVEN-G_Metodologia_fundacional.html#8-1-use-case-roles
- **AI risk owner**: A second-line function that reviews the risks, regulatory classification and suppliers of each use case, and confirms the intensity. It is independent from whoever builds: it cannot verify what it has prepared. It prevents the team that wants to launch the use case from assessing its own risks alone. The link opens the role incompatibilities in document 01. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/01_SEVEN-G_Metodologia_fundacional.html#8-2-incompatibilities
- **Application intensity**: How much governance each use case needs: Express, Lite or Enterprise. It is decided per use case, not per company. The gates are always the same; what changes is how many criteria are assessed, which sessions are grouped and who verifies. It makes the framework proportionate: an internal assistant does not carry the same load as an agent deciding about customers. The link opens the definition in document 01. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/01_SEVEN-G_Metodologia_fundacional.html#application-intensity-express-lite-and-enterprise

SOURCES AND LINKS

- **Document 33 §4.3 · Risk level and 5 × 5 matrix** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#4-3-risk-level-and-5-5-matrix
- **Document 33 §5.3 · Control effectiveness** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#5-3-control-effectiveness
- **Document 32 §5.4 · Illustrative classification examples** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#5-4-illustrative-examples
- **Risk assessment scales** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#assessment-scales

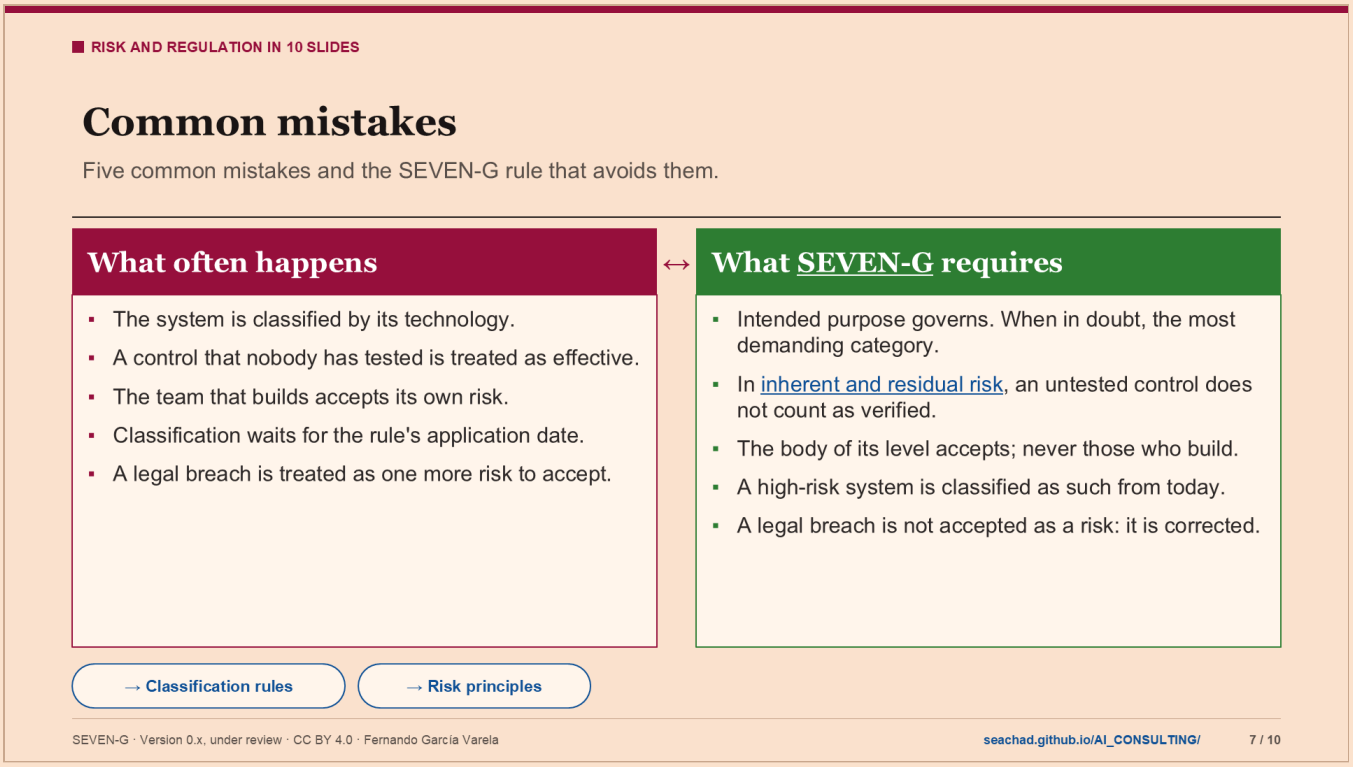

Notes

Each mistake on the left has a rule written in documents 32 and 33 that avoids it.

- Classifying by technology. The same language model can be used to summarise internal documents or to filter applications in a recruitment process. The first is usually minimal risk; the second, high risk. That is why the intended purpose and reasonably foreseeable use govern. When in doubt, the most demanding category applies while it is resolved, with a deadline.
- Treating an untested control as effective. A written policy or an instruction to the model is not, on its own, an effective control against prompt injection: it must be tested with adversarial tests. An untested control can be used in the target residual, but not in the verified one.
- The team accepting its own risk. This is a conflict of interest. SEVEN-G segregates duties: no one accepts a risk of a use case they are building.
- Waiting for the application date. According to document 34, at the consultation date, the obligations for Annex III high-risk systems apply from 2-12-2027. SEVEN-G asks for them to be classified as such from today, so that the design builds in the requirements in time and systems do not have to be rebuilt in a hurry.
- Accepting a legal breach. Acceptance is for risks within the appetite. A breach of the law is not a risk to accept: it is corrected.

There is a sixth mistake worth discussing in the classroom: not looking for the systems nobody has declared. Document 32 asks for the inventory to be reconciled, at least every quarter, with independent sources: procurement, cloud invoices, code repositories and expense claims.

Question for the classroom: which of these mistakes have you seen in your organisation, and what made it possible?

TERMS ON THIS SLIDE

- **SEVEN-G**: SEVEN-G (Seven-phase Enterprise Value & Governance) is an open framework for a company to implement, govern and measure AI, and to know whether it is transforming or only gaining efficiency. It treats each use case as an investment with decision gates and gives the board figures it can trust. Management, the AI office, risk and the board use it. The link opens its entry page, with the story, the questions and the examples. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/entrada/index.html
- **Inherent and residual risk**: Inherent risk is the risk a use case has without controls; residual risk is what remains after applying them. Both are measured as likelihood by impact on a 5 × 5 matrix. The residual level decides who must accept the risk: only the board accepts a Critical one. The link opens the explanation in document 33. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#inherent-risk-controls-and-residual-risk

SOURCES AND LINKS

- **Document 32 §5.3 · Classification application rules** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#5-3-application-rules
- **Document 32 §9 · Detection of uninventoried systems** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#detection-of-uninventoried-systems
- **Document 33 §5.3 · Control effectiveness** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#5-3-control-effectiveness
- **Document 34 §3.1 · Verified timeline of application** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html#3-1-verified-timeline-of-application
- **Regulation (EU) 2026/1744 (digital omnibus on AI), EUR-Lex** https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng
- **Classification rules** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#5-3-application-rules
- **Risk principles** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#principles

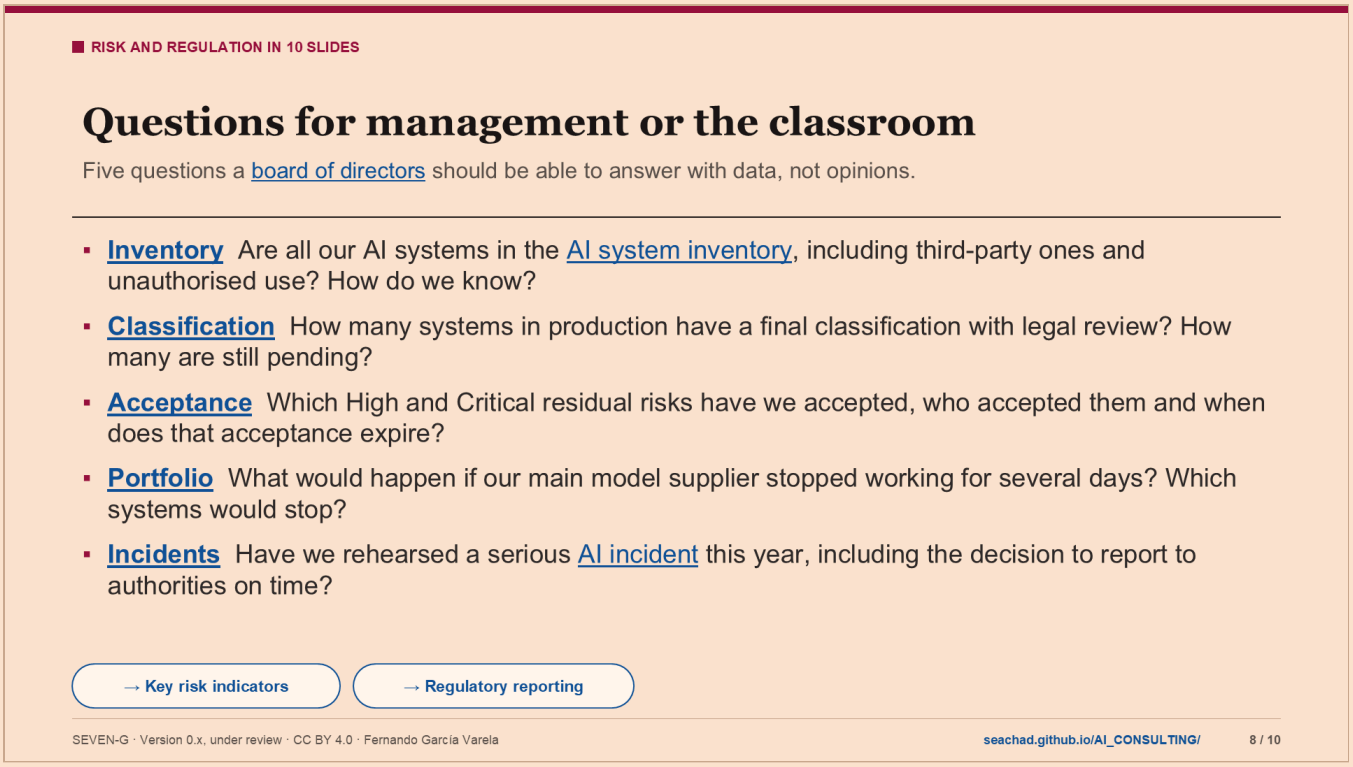

Notes

These questions are useful for a board or committee meeting, and also for a classroom debate with a real or fictitious case. Each one has an indicator or a rule of the framework behind it.

- Inventory. Document 32 measures inventory coverage: the systems that appear in sources outside the inventory and were not registered. Its indicative target is to stay below 5 %, with a decreasing trend.
- Classification. The indicator “final classification with legal review” has a target of 100 % of systems in production. A system in production without a complete record and a final classification is a major nonconformity; if it is also high risk or prohibited, a critical one.
- Acceptance. The key risk indicators in document 33 include Critical residuals (zero without board approval), High residuals and expired acceptances (zero).
- Portfolio. The risks of use cases do not add up: they interact. Document 33 recommends portfolio scenarios, such as the unavailability of the main model supplier for several days or a behaviour change in a base model used by several systems.
- Incidents. The same incident can trigger several reporting regimes at once. Document 37 summarises them, for example the EU AI Act regime for serious incidents or the GDPR regime for personal data breaches, and asks for reports to be rehearsed at least once a year in an S1 incident drill. Before applying them, each organisation verifies the time limits in force with its legal advisers.

The indicator thresholds are indicative. Each company approves them at the direction stage (C2) and recalibrates them at the annual review (C5).

Classroom exercise: choose one of the five questions and say what data you would need to answer it, who holds it and how often it should reach the board.

TERMS ON THIS SLIDE

- **Board of directors**: The company's highest body. In SEVEN-G it approves the AI thesis and risk appetite (C2), Transform bets and their scaling, and oversees AI every quarter. AI is a strategic asset: its big decisions are not only

for technology. The link opens its role in the governance model (document 30). https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/30_SEVEN-G_Modelo_de_gobierno.html#3-1-board-of-directors

- **AI system inventory** (T02): A list of all the company's AI systems: its own, third-party AI embedded in processes, the assistants employees use and any unauthorised use detected. Each system has an owner, a regulatory classification, an intensity and an autonomy level. It comes first, because you cannot govern what you do not know. It is a view of the use case register: open the link and choose «Inventory» (fictitious data). https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html
- **AI incident and its severity** (S1–S4): An event in which an AI system causes or may cause harm: serious error, bias, data leak, improper action by an agent. It is classed from S1 (the most serious) to S4. S1 and S2 reach the board by the fast track and may require notifying authorities (P51). The link opens the severity scale in document 37. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#4-2-severity

SOURCES AND LINKS

- **Document 32 §8.1 · Inventory quality indicators** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#8-1-indicators
- **Document 33 §10 · Portfolio risk** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#portfolio-risk
- **Document 33 §11.2 · Key risk indicators** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#11-2-key-risk-indicators
- **Document 37 §5 · Regulatory reporting** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#regulatory-reporting
- **Regulation (EU) 2016/679 (GDPR), EUR-Lex** https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng
- **Key risk indicators** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#11-2-key-risk-indicators
- **Regulatory reporting** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#regulatory-reporting

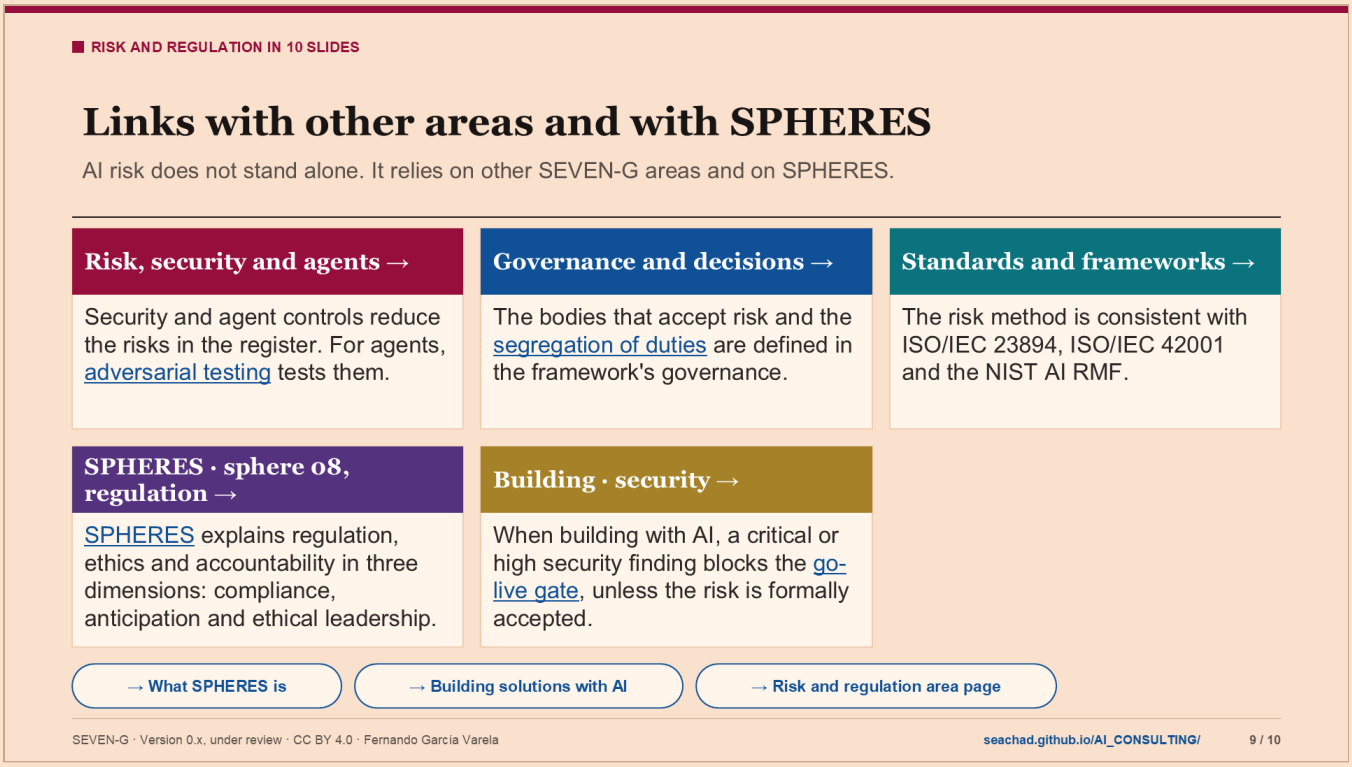

Notes

The risk and regulation area is connected with other SEVEN-G areas and with SPHERES.

- Risk, security and agents. Document 35 defines security and agent controls, with codes. At the design gate (G4), every Medium or higher risk has a designed control, with its code where one exists. In generative AI and agents, effectiveness against prompt injection must be tested with adversarial tests.
- Governance and decisions. The acceptance scale uses the framework's bodies: product owner, sponsor, AI Committee and board. The rule that those who build do not accept their own risk comes from segregation of duties.
- Standards and frameworks. The risk methodology is consistent with ISO/IEC 23894, with the risk requirements of ISO/IEC 42001 and with the Map, Measure and Manage functions of the NIST AI RMF. It does not replace the company's corporate risk framework: it specialises it for AI and is integrated into it.
- SPHERES, sphere 08. SPHERES is the explanatory layer of the nine spheres. Sphere 08, Regulation, ethics and accountability, looks at three dimensions: compliance (today), anticipation (tomorrow) and ethical leadership (beyond the law). Rule for the map: sphere 08 is never used to signal that a use case in another sphere has regulatory risk; that goes into its classification and its risk register.
- Building the solution, security. SEVEN-G document 53 requires security, secrets and dependency analysis on AI-generated code before the go-live gate (G5). A critical or high finding blocks G5, unless the risk is formally accepted using the scale of document 33. For critical security controls, continuing with conditions is not allowed.

Why it matters: the same risk is seen from the business (SPHERES), from governance and measurement (SEVEN-G) and from building the solution (document 53), with the same language.

Question for the classroom: who should be able to veto a legal use of AI that the company considers unacceptable?

TERMS ON THIS SLIDE

- **Adversarial testing** (P53): Tests in which a team tries to trick or break an AI system as an attacker would: prompt injection, data leakage, abuse of an agent's tools. A critical scenario in which the attacker succeeds prevents deployment. They are planned and reported with template P53. The link opens the template. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P53_SEVEN-G_Plan_e_informe_de_pruebas_adversarias.html
- **Segregation of duties**: The principle whereby whoever builds does not control and nobody approves their own work: whoever provides the evidence neither verifies it nor decides, and incompatible roles do not fall to the same person. It is one of the core rules that are never lowered. It assures the board that a figure or a control is not signed only by whoever defends it. The link opens the roles and incompatibilities in document 01. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/01_SEVEN-G_Metodologia_fundacional.html#roles-bodies-and-segregation-of-duties
- **SPHERES**: A methodology that explains where AI can create value in a company and with what ambition, through three questions: where do we play?, with what ambition? and what do we ask? It is the explanatory layer of the spheres and ambition levels; the operating rules are in SEVEN-G. It supports conversations with management and the board. The link opens its presentation document. https://seachad.github.io/AI_CONSULTING/SPHERES/html/en/00_SPHERES_Que_es_y_para_que_sirve.html#what-spheres-is
- **Go-live gate** (G5): The gate that authorises using the AI system with real users, customers or data. It requires validation and pilot passed, a tested rollback plan and multi-level sign-off (business, technology, risk). Nothing goes into production without this sign-off. The link opens its criteria in document 21. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/21_SEVEN-G_Criterios_de_gate_y_auditoria.html#6-6-g5-go-live

SOURCES AND LINKS

- **SPHERES 04 · Sphere 08, Regulation, ethics and accountability** https://seachad.github.io/AI_CONSULTING/SPHERES/html/en/04_SPHERES_Regulacion_y_gobierno_de_la_IA.html#sphere-08-regulation-ethics-and-accountability
- **Document 53, section 7.4 · Security** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/53_SEVEN-G_Construccion_de_soluciones_con_IA.html#7-4-security
- **Document 33 §1.3 · Reference standards** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#1-3-reference-standards
- **ISO/IEC 23894:2023 · Guidance on AI risk management** https://www.iso.org/standard/77304.html
- **NIST AI RMF 1.0 · AI Risk Management Framework** https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
- **What SPHERES is** https://seachad.github.io/AI_CONSULTING/SPHERES/html/en/00_SPHERES_Que_es_y_para_que_sirve.html
- **Building solutions with AI** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/53_SEVEN-G_Construccion_de_soluciones_con_IA.html
- **Risk and regulation area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/riesgo.html

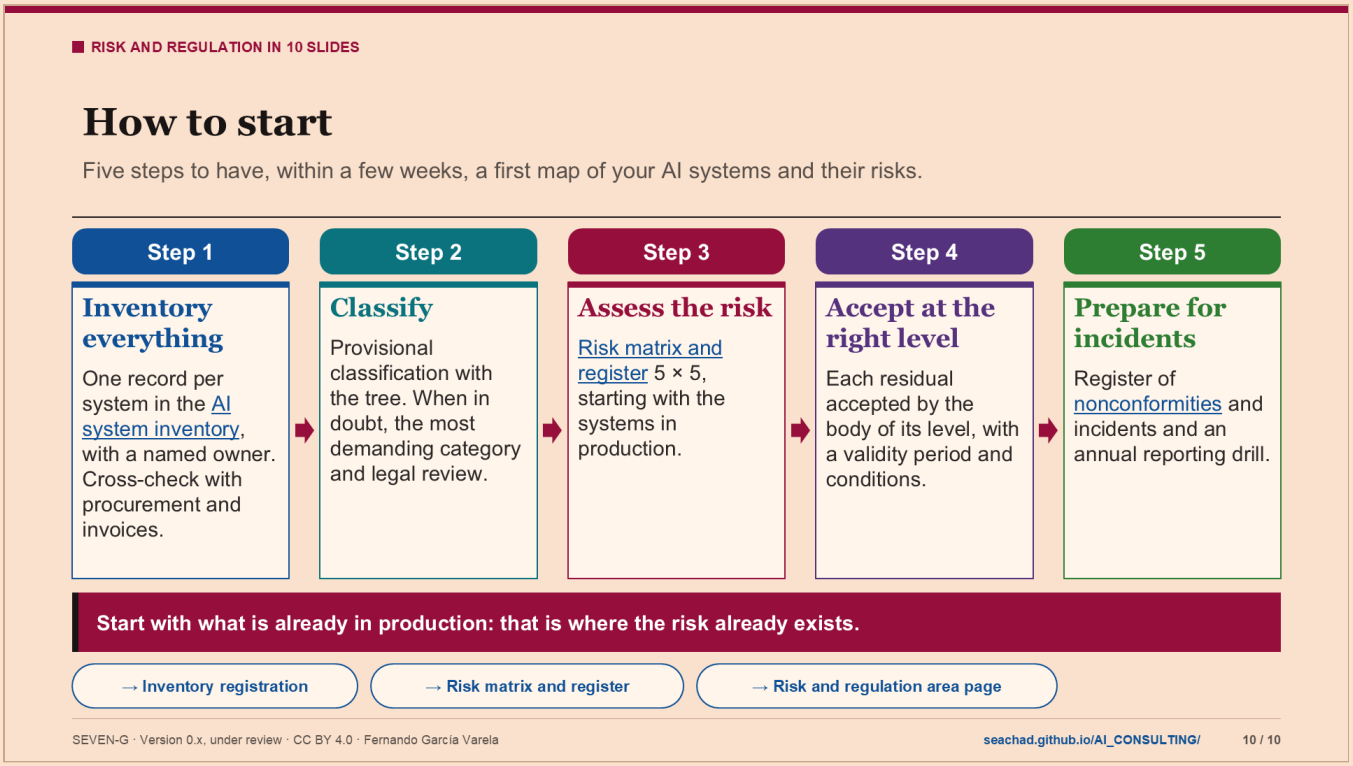

Notes

These five steps put in order what documents 32, 33 and 37 ask of a company that is starting.

- Step 1. Inventory everything. Document 32 gives reference time limits: registration within 5 working days of identifying the system; provisional classification within 10 working days of registration. Each system has a named owner, not an area. To find what nobody has declared, the inventory is reconciled with procurement and contracts, cloud invoices, code repositories and expense claims. Template: inventory record (P05).
- Step 2. Classify. The classification tree goes through seven questions: whether it is an AI system, whether an exclusion applies, whether it is a prohibited practice, whether it is high risk under Annex I or Annex III, whether the Article 6(3) exception applies and whether there are transparency obligations. Then the company's role and the resulting assessments are set. Template: P11. The final classification requires qualified legal judgement.
- Step 3. Assess the risk. Start with the systems in production, which are the ones that can already cause harm. Each risk is described as cause, event and consequence, assessed with the 5 × 5 matrix and given a named owner. Template: P12; tool: the risks view of the register (T06).
- Step 4. Accept at the right level. Each residual risk is accepted by the body of its level, with a validity period and conditions. A legal breach is not accepted: it is corrected.
- Step 5. Prepare for incidents. Open the nonconformity and incident register (P50) and rehearse reports at least once a year in an S1 incident drill (P51).

The time limits and thresholds are for reference: each company sets them at the direction stage (C2). And, once again, SEVEN-G is an aid, not legal advice: each organisation verifies the regulation that applies to it.

To continue, open the area page. From there you reach documents 32, 33, 34 and 37 and template P12.

TERMS ON THIS SLIDE

- **AI system inventory** (T02): A list of all the company's AI systems: its own, third-party AI embedded in processes, the assistants employees use and any unauthorised use detected. Each system has an owner, a regulatory classification, an intensity and an autonomy level. It comes first, because you cannot govern what you do not know. It is a view of the use case register: open the link and choose «Inventory» (fictitious data). https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html
- **Risk matrix and register** (T06): A likelihood-by-impact matrix (5 × 5) and register of each use case's risks: inherent and residual risk, controls and their effectiveness, response, contingency and who accepts the risk. It is mandatory from phase 3. It is filled in with template P12 or in the «Risks» view of the use case register, which also downloads an Excel file for the use case with proposed mitigation plans. Fictitious data. https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T01_registro_iniciativas/registro.html
- **Nonconformity** (P50): A breach of a framework rule or a control, even if it caused no harm (for example, a use case in production without sign-off). It is classed as minor, major or critical, recorded (P50), its root cause analysed (P52) and corrected on time. It lets governance learn from its failures. The link opens the process in document 37. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#nonconformity-process

SOURCES AND LINKS

- **Document 32 §7.1 · Inventory registration procedure** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#7-1-registration
- **Document 32 §5 · Regulatory classification tree** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#regulatory-classification-tree
- **Document 33 §8 · Risk register** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/33_SEVEN-G_Metodologia_de_riesgos_de_IA.html#risk-register
- **Document 37 §3.10 · Nonconformity time limits** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/37_SEVEN-G_No_conformidades_e_incidentes.html#3-10-time-limits
- **Risk and regulation area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/riesgo.html
- **Inventory registration** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/32_SEVEN-G_Inventario_y_clasificacion_regulatoria.html#7-1-registration
- **Risk matrix and register** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P12_SEVEN-G_Matriz_y_registro_de_riesgos.html

LEGAL NOTICE

Legal notice: a generic, free methodological aid provided "as is" for information purposes; it is not legal advice and has not been legally reviewed for any case. Each organisation adapts it and is responsible for its legal review and its regulatory compliance; it may not reflect recent regulatory changes. To the fullest extent permitted by law, the author accepts no liability whatsoever for the effects of its application. It does not grant certification of any kind. Details in document 93.

CC BY 4.0 (content) · MIT (code) · Fernando García Varela · seachad.github.io/AI_CONSULTING/