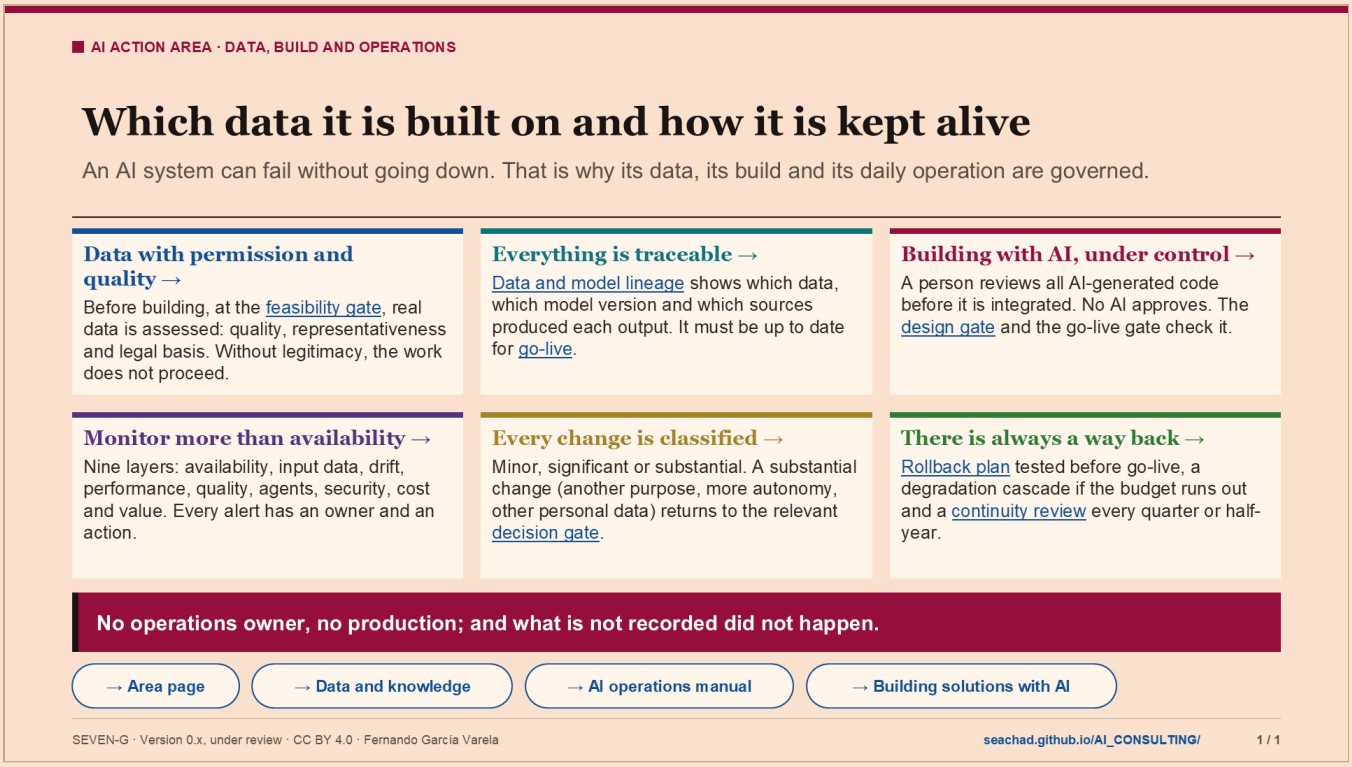


Notes

This slide summarises the Data, build and operations action area of SEVEN-G. It answers a question that boards rarely ask, but which decides whether AI brings lasting value: which data an AI system is built on and how it is kept alive in production. An AI system can stay available and yet give worse answers, treat a group of people differently or spend twice as much. That is why SEVEN-G governs three moments: the data before building, the build and daily operation.

The six blocks, one by one:

- Data with permission and quality. In phase 3 (feasibility) data is assessed with real data or a sample from the source systems, not with descriptions. Quality, representativeness, whether the data will be the same in production and whether there is a legal basis to use it are all measured. The data owner, who sits in the business, authorises the use at the feasibility gate (G3). If the legitimacy block is insufficient, the use case does not proceed. The reason: a data problem discovered in phase 5 costs the entire build.
- Everything is traceable. Data and model lineage (P16) makes it possible to reconstruct which data, which transformations, which model version and which sources produced an output. It is the basis for investigating an incident, responding to an affected person or auditing. It is prepared in phase 4 and must be up to date for go-live (G5).
- Building with AI, under control. More and more code is written with AI assistants. SEVEN-G requires a competent person to review it before it is integrated, that it is known which part was generated with AI, that tests are defined before implementation and that critical or high security findings block G5. No AI approves: every acceptance is recorded by an identified person. The design gate (G4) checks that the build tools are approved and that the test strategy exists; the go-live gate (G5), that the human review is recorded and that no critical or high findings remain open.
- Monitor more than availability. The operations manual (document 52) defines nine layers: availability, input data quality, drift, model performance, response quality, agent actions, security, cost, and human

oversight with value. Drift includes usage drift: the assistant starts to receive questions that nobody validated.

- Every change is classified. A minor change does not alter behaviour. A significant one alters it within what was approved and requires regression evaluation. A substantial one alters the purpose, scope, autonomy, data, critical controls or risk, and requires a return to the relevant gate. When in doubt, the higher class applies.
- There is always a way back. The rollback plan (P19) is tested before G5 and then periodically. If the consumption budget runs out, a degradation cascade reduces cost without touching the controls. The continuity review (R6) replaces the gate in phase 6: quarterly in Enterprise and half-yearly in Lite.

Why it matters: most of the value of AI is won or lost after go-live. A board that only approves projects and does not ask how they are operated sees the spending, but not the degradation.

Question for the classroom: who in your organisation is today the reachable person accountable for an AI system in production when something fails on a Sunday?

To go further: the area page, documents 51 (data and knowledge), 52 (operations manual) and 53 (building with AI), and its section 8, which says what is checked at each gate.

TERMS ON THIS SLIDE

- **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
- **Data and model lineage** (P16): A template that records where the data comes from, which models and versions are used, how they were trained or configured and which fallback model exists. It makes it possible to reconstruct what happened if something fails and to change model without surprises. The link opens the template. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P16_SEVEN-G_Linaje_de_datos_y_modelos.html
- **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
- **Design gate** (G4): The gate that approves the full design before building: architecture, data, security, human oversight, rollback and adoption. It checks that the controls are planned from the start. 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#6-5-g4-design
- **Decision gates** (G0–G7): At the end of each phase there is a gate: a moment when a body decides, with verified evidence, whether the use case continues, continues with conditions, iterates, pivots or stops. There are at most two iterations before the decision is escalated. Gates turn enthusiasm into recorded decisions. The link opens their definition in document 01. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/01_SEVEN-G_Metodologia_fundacional.html#decision-gates
- **Rollback plan** (P19): A tested plan to go back if the AI system fails: how it is switched off, which non-AI alternative is used and how fast. For agents it includes a kill switch. It is designed before go-live and actually tested (template P19). The link opens the template. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P19_SEVEN-G_Plan_de_reversion_y_prueba.html
- **Continuity review** (R6): A periodic review of a use case in production: quarterly in Enterprise, half-yearly in Lite and yearly in Express. It checks whether value is still arriving and whether there is drift, incidents or regulatory change. A deviation brings forward the scale or retire decision (G7). It is prepared with pack P65. 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#6-7-r6-continuity-review

SOURCES AND LINKS

- **Data, build and operations area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/operacion.html
- **Document 51 · Data and knowledge** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/51_SEVEN-G_Datos_y_conocimiento.html
- **Document 52 · AI operations manual** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html
- **Document 53 · Building solutions with AI** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/53_SEVEN-G_Construccion_de_soluciones_con_IA.html
- **Document 53, section 8 · Verification at the gates** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/53_SEVEN-G_Construccion_de_soluciones_con_IA.html#verification-at-the-gates
- **Glossary (document 02)** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/02_SEVEN-G_Glosario_y_convenciones.html#glossary-of-framework-terms
- **Area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/operacion.html
- **Data and knowledge** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/51_SEVEN-G_Datos_y_conocimiento.html
- **AI operations manual** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html
- **Building solutions with AI** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/53_SEVEN-G_Construccion_de_soluciones_con_IA.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/