TEMPLATE P25 · MONITORING AND ALERTING CONFIGURATION

Monitoring and alerting configuration

Defines what is monitored in the AI system in production, with which thresholds, who receives each alert and what action it triggers.

Fernando García Varela Version 0.1 (working draft) 16-09-2026 Draft for review

IN THIS DOCUMENT

01 Instructions for use

02 Identification

03 Metrics catalogue

04 Minimum metrics by type of system

05 Alert rules

06 Dashboards and recipients

07 Configuration testing

08 Threshold review

09 Quality criteria

10 Approval and verification

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regulation specific to each sector or jurisdiction may be incomplete, may not apply to a particular case or may become out of date as a result of regulatory changes, interpretations or supervisory positions after their consultation date.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

01

Instructions for use

Aspect	Guidance
When it is used	It is designed in phase 4, configured and tested in phase 5 and maintained in phase 6.
Who completes it	AI Operations Owner, with the AI Technical Owner and the AI Product Owner (adoption and value metrics).
Who verifies it	AI Office at Lite intensity; AI Auditor at Enterprise intensity.
At which *gate* it is reviewed	G4 (design), G5 (tested configuration) and R6 (currency of thresholds).

Aspect	Guidance
Lite versus Enterprise	Lite may omit the fields marked **(Enterprise)**. Every use case in production must monitor, as a minimum, performance, availability, cost and the human oversight indicators of P17.
Relationship	P09 Baseline · P16 Lineage (drift reference) · P17 Oversight indicators · P18 Action log · P19 Rollback triggers · P24 Operations manual · P26 Incident response · P28 Value tracking · documents 41 (Indicator catalogue) and 52.

Rules:

- Each rollback trigger in P19 must have an alert in this document.
- **"No data" is not zero:** the absence of a metric is in itself an alert, never a normal value.
- The severity of an alert is an **initial proposal**; the severity of the incident is classified using the criteria of document 37 (P26).
- Threshold changes are recorded in P27 with their reason.

02

Identification

Field	Content	Guidance
Use case code		IA-AAAA-NNN (T01).
AI system(s)		Identifier in the inventory (T02).
Technology		Predictive ML · Generative AI · Agent · Language and document processing · Vision · Optimisation · Embedded third-party AI · Rules (not AI).
Intensity		Lite · Enterprise.
Monitoring tool		Name of the internal or provider tool.
Document version		
Author		Role and name.
Date		DD-MM-YYYY.
Document status		Draft · Submitted for verification · Verified · Approved · Superseded.

03

Metrics catalogue

Category: Model performance · Data drift · Usage drift · Input data quality · Bias and fairness · Availability and latency · Cost and consumption · Security · Agent behaviour · Human oversight · Adoption · Value.

ID	Metric	Category	Definition and formula	Source	Calculation frequency	Reference value	Warning threshold	Critical thresh
M-01 ***(illustrative example)***	Human override rate	Human oversight	Proposals rejected or corrected ÷ proposals reviewed, daily	Validation log	Daily	8% (pilot)	> 15%	> 25% o two consecu days
M-02 ***(illustrative example)***	Model consumption cost	Cost and consumption	Daily sum invoiced by the provider	Provider console	Daily	€65/day	> €90/day	> €150/
M-03 ***(illustrative example)***	Actions blocked by limits	Agent behaviour	No. of actions blocked by limits or intent-based access control	Action log (P18)	Continuous	0–3/day	> 10/day	Any acti execute outside limits
M-04 ***(illustrative example)***	Input data freshness	Input data quality	Hours since the last successful load	Load process	Hourly	< 2 h	> 4 h	> 12 h or data

04

Minimum metrics by type of system

Tick those that apply and link the ID from block 3.

Metric	Predictive ML	Generative AI	Agent	ID in block 3
Performance against reference (with actual outcome when available)	☐	☐	☐	
Drift in input variables	☐			
Usage drift: distribution of topics or intents and queries outside the validated scope		☐	☐	
Performance differences between groups (Enterprise)	☐	☐	☐	
Unequal responses in counterfactual pairs (mandatory if the system decides, recommends or communicates with people, or has direct exposure)		☐	☐	
Ungrounded responses or responses flagged as incorrect by users		☐	☐	
Detection of prompt injection attempts		☐	☐	
Detected leakage of confidential or personal data		☐	☐	
Actions executed, blocked and validated by people			☐	
Kill switch activations			☐	
Human override rate and validation time	☐	☐	☐	
Availability, latency and errors	☐	☐	☐	
Consumption and infrastructure cost	☐	☐	☐	
Days in cost-driven degraded mode and quality during the mode		☐	☐	
Effective use (adoption)	☐	☐	☐	

References for generative AI and agents (document 52 §4.2.6, §4.2.7 and §10.3):

Field	Content	Guidance
Usage reference		Topics, intents and scope validated at G5, with their reference distribution (P16).
Query classification		Rules, classifier or human sampling; sample size and frequency.

Field	Content	Guidance
Counterfactual pairs		Protected attributes or proxies tested (according to the impact assessment), number of pairs and what counts as a material difference (decision, amount, tone, refusal, completeness). Threshold in P17.
Minimum quality in degraded mode		Minimum value of the quality metric below which the system moves to the next level of the cascade; sampling at least doubled during the mode.

05

Alert rules

Proposed initial severity: S1 Critical · S2 High · S3 Medium · S4 Low (confirmed using document 37).
Associated action: Review · Apply P24 procedure · Open incident (P26) · Activate rollback trigger (P19) · Kill switch.

Alert ID	Condition (metric and threshold)	Initial severity	Recipient	Channel	Maximum acknowledgement time	Associated action	P19 trigger
AL-01 ***(illustrative example)***	M-01 > 25% on two consecutive days	S2 High	AI Operations Owner	On-call tool	2 working hours	Activate rollback trigger	D-01
AL-02 ***(illustrative example)***	M-03: action executed outside limits	S1 Critical	Technical on-call and security	Immediate notification	15 minutes	Kill switch and open incident	D-02
AL-03 ***(illustrative example)***	Any critical metric with no data for more than one cycle	S3 Medium	Technical on-call	On-call tool	4 working hours	Apply P24 procedure	—
AL-04 ***(illustrative example)***	Month-to-date cost ≥ 100% of the budget	S4 Low	AI Operations Owner	On-call tool	4 working hours	Apply P24 procedure (cost-driven degradation cascade, level N2)	D-04

06

Dashboards and recipients

Dashboard or report	Recipients	Content	Frequency
Operational dashboard	Operations owner and on-call	All block 3 metrics in near real time	Continuous
Oversight report	Product owner and supervisors	Human oversight indicators and complaints	Weekly
Report for R6	R6 decision-maker	Trends, alerts, incidents and value	According to intensity
Data for the board dashboard (Enterprise)	AI Office	Aggregated indicators in the dashboard format (T17)	Quarterly

07

Configuration testing

Each alert with a rollback, kill switch or incident action is tested before G5.

Alert ID	Test date	Method (Simulation of the condition · Injection of test data · Other)	Did it reach the recipient on time? (Yes · No)	Result (Passed · Not passed)	Observations

08

Threshold review

Field	Content	Guidance
Review frequency		At least at each R6.
Who proposes changes		AI Operations Owner.
Who approves changes		Technical owner with risk clearance for critical thresholds.

Field	Content	Guidance
Threshold change log		P27, with reason and previous value.
Recalibration after relevant changes (Enterprise)		A new model version or population entails reviewing reference values.

09

Quality criteria

Formal G4, G5 and R6 criteria in document 21.

#	Check	Status (Met · Not met · Not applicable · Pending)
1	Each metric has a definition, source, reference, thresholds and owner.	
2	The minimum metrics for the type of system are covered.	
3	Every P19 trigger and every P17 oversight indicator has an alert.	
4	The absence of data generates an alert.	
5	Critical alerts have been tested and reach the recipient on time.	
6	Threshold changes are recorded with their reason.	
7	For generative AI and agents, the usage reference, the counterfactual pairs where applicable and the minimum quality in degraded mode are set.	

10

Approval and verification

Function	Role	Name	Date	Signature or reference
Author	AI Operations Owner			
Co-author	AI Technical Owner			
Clearance of critical thresholds	AI Risk Owner			

Function	Role	Name	Date	Signature or reference
Verifier	AI Office (Lite) · AI Auditor (Enterprise)			
***Gate* decision-maker**	As per 01 §7.5 (recorded in P29)			

Segregation of duties: the verifier is not part of the team that builds or operates the system.

11

Version control

Version	Date	Changes
0.1	16-09-2026	First version of the template.
0.1	18-09-2026	Usage drift, unequal responses in counterfactual pairs and days in cost-driven degraded mode in the minimum metrics; references for generative AI and agents; example alert AL-04; quality criterion 7.