



AGENTIC COMMERCE FRAMEWORK®

THE MANUAL

Governing Autonomy

The ACF® toolkit gathered in a single volume

Forty-five operational cards in three families (decide, organise, operate), fifteen doctrinal markers, two running cases that walk through the framework from A to Z and a complete user guide, to **frame, deploy, supervise and audit** AI agents in your organization, from initial diagnostic to day-to-day operation.

45

A4 CARDS

3

FAMILIES · DECIDE ·
ORGANISE · OPERATE

Free

SCHOOLS & TRAINING

Vincent Dorange · 2026 Edition




The standard, the tools & the educator space

Scan · acfstandard.com

The ACF pyramid

The image that captures the entire framework: **delegate autonomy at the apex** without ever loosening the **foundation of sovereignty** that supports it. The higher you delegate, the broader the governance base must be.



TRACEABLE ACCOUNTABILITY

The ground: every decision remains **attributable to a named human**, the DDAO · register **ACF-O9**, exceptions **ACF-O8**, redress **ACF-O10**, records **ACF-P8 · P13 · P14**

Delegated autonomy

APEX

The agent acts · narrow perimeter, traceable, reversible.

ACF-05

ACF-08

ACF-P1

ACF-P2

ACF-P3

Ultimate human control

PRINCIPLE

The hand never relinquished: supervise, take over, stop.

ACF-06

ACF-07

ACF-10

ACF-11

ACF-12

ACF-16

ACF-17

ACF-O5

ACF-O6

ACF-P4

ACF-P6

ACF-P7

Governance by design

PRINCIPLE

The rules before deployment: mandates, thresholds, perimeters.

ACF-02

ACF-03

ACF-04

ACF-13

ACF-14

ACF-15

ACF-O4

ACF-O9

ACF-P5

ACF-P9

ACF-P15

Decisional sovereignty

FOUNDATION

The foundation: mastery of your decisions, data, and dependencies.

ACF-01

ACF-O1

ACF-O2

ACF-O3

ACF-O7

ACF-P10

ACF-P11

ACF-P12

Reading · the apex (**autonomy**) only holds on the width of the foundation (**sovereignty**); the entire structure rests on **traceability**.

Table of Contents

INTRODUCTION & DOCTRINE

-	The problem: autonomy without governance	
-	What is ACF?	<i>the standard in brief</i>
-	The third European way	
-	The 4 new laws of robotics	
-	The 4 founding principles	
-	The 4 layers of governance	
-	The 4 levels of delegated autonomy	
ACF-03	The criticality score	
ACF-12	The risk matrix	
-	The 4 levels of agentic maturity	
-	The role of the DDAO	
-	The 6-step journey	

THE THREE FAMILIES

Decide	ACF-01 → ACF-20	20 fiches
<i>Who authorised this decision, and how far does it go?</i>		
Organise	ACF-01 → ACF-O10	10 fiches
<i>Who actually decides, now?</i>		
Operate	ACF-P1 → ACF-P15	15 fiches
<i>Has the signed rule become verifiable behaviour?</i>		

Every card comes three times: the blank form, then the same form filled in by Lendari, then by Rhena.

USER GUIDE & APPENDICES

-	Toolkit user guide	<i>where to start</i>
-	Frequently asked questions (FAQ)	<i>20 answers</i>
-	Resources & references	
-	The ACF tools	<i>ecosystem</i>
-	Conclusion	

WHERE TO START



Then follow the **6-step journey** · detailed in the doctrine.

THIS DOCUMENT IS PROTECTED

The toolkit is free for teaching, it is not free of rights. A teacher may pass it whole to their students, unmodified, on a closed course space. Reproducing part of it, translating it, publishing it under another name, embedding it in a commercial offering or using it to train a model requires written permission. The full terms close this manual and are published at acfstandard.com/licence.

The 45 cards

Decide · Organise · Operate

acfstandard.com

DECIDE · ACF-01 TO ACF-20

ACF-01	Sovereignty Score
ACF-02	Decision Map
ACF-03	Criticality Matrix
ACF-04	Agentic Constitution
ACF-05	Agent Card
ACF-06	Supervision & Governance
ACF-07	Kill switch
ACF-08	First Agent Dossier
ACF-09	Agentic Decisions Register
ACF-10	Action & Improvement Plan
ACF-11	30-Day Governance Audit
ACF-12	Agentic Risk Assessment
ACF-13	Agent Mandate
ACF-14	Responsibility by Design
ACF-15	Objective and Counter-Metric
ACF-16	Supervision Test
ACF-17	Validity Review
ACF-18	Guided Case Study
ACF-19	Instructor's Guide
ACF-20	Governance Simulation

ORGANISE · ACF-O

ACF-O1	Human-AI-Agent Matrix
ACF-O2	Hybrid Organisation Map
ACF-O3	Critical Skills to Keep
ACF-O4	AI Policy
ACF-O5	Shadow AI, undeclared usage
ACF-O6	Shadow Agents, undeclared agents
ACF-O7	The Shift of Power
ACF-O8	The Exceptions Register
ACF-O9	Governing the Fleet
ACF-O10	Affected People and Their Appeal

OPERATE · ACF-P

ACF-P1	Agent Configuration Grid
ACF-P2	Acceptance Protocol
ACF-P3	Go-Live Ramp
ACF-P4	Weekly Supervision Log
ACF-P5	Role Allocation
ACF-P6	Escalation & Confidence Card
ACF-P7	Stop Mechanism Test
ACF-P8	Incident Log
ACF-P9	Permissions & Data Review
ACF-P10	Vendor Assessment Grid
ACF-P11	Foundation Dependency Map
ACF-P12	Degraded Mode & Exit Plan
ACF-P13	Model Version Log
ACF-P14	Holding the Record
ACF-P15	Blind Spot Statement

THE TWO RUNNING CASES

Lendari

Online fashion brand, €120M revenue, 90 staff, Paris. Agent "Margaux", level 1 customer service and returns.

Rhena

German insurer, €75bn in premiums, 62,000 staff, 28 countries, Cologne. Agent "ANNA", level 1 home claims.

Every card is read three times: the blank form, then the same form filled in by each of the two cases, page for page. **Other sectors are covered in the sector handbooks**, published separately: health, public administration, banking, industry, transport, energy, retail, telecommunications, education, etc.

Where to start

You do not need all 45 cards to govern one agent. The toolkit is modular: take the path that matches your situation, and open the other cards only on the day the question arises. Four paths start from a situation, two run a whole family.

11 CARDS · ONE DAY

I am launching my first agent

You are starting from nothing. The point is not to cover everything, it is to delegate nothing you have not written down.

ACF-01 · ACF-02 · ACF-03 · ACF-05 · ACF-08 · ACF-13 · ACF-06 · ACF-07 · ACF-01 · ACF-P1 · ACF-P6

9 CARDS · ONE DAY

I am auditing an agent already live

The agent is running. You are looking for what the setup does not see, and you want to be able to prove it.

ACF-09 · ACF-11 · ACF-12 · ACF-16 · ACF-17 · ACF-P4 · ACF-P8 · ACF-P9 · ACF-P14

8 CARDS · TWO WORKSHOPS

I am preparing my organisation

No agent is live yet, or several already are without anyone having decided so.

ACF-04 · ACF-15 · ACF-01 · ACF-02 · ACF-03 · ACF-04 · ACF-05 · ACF-07

9 CARDS · QUARTERLY REVIEW

I am governing a fleet

The first agent is governed by attention. The tenth is governed by inventory, or it is not governed at all.

ACF-06 · ACF-08 · ACF-09 · ACF-010 · ACF-P10 · ACF-P11 · ACF-P12 · ACF-P13 · ACF-P15

10 CARDS · THE ORGANISE FAMILY

I am bringing my organisation up to speed

The agent is never alone: it shifts roles, skills and power. This family runs end to end, in order.

ACF-01 · ACF-02 · ACF-03 · ACF-04 · ACF-05 · ACF-06 · ACF-07 · ACF-08 · ACF-09 · ACF-010

15 CARDS · THE OPERATE FAMILY

I am running the agent day to day


From the setting to the proof. This family runs end to end, in the order of an agent's life, from configuration to leaving the vendor.

ACF-P1 · ACF-P2 · ACF-P3 · ACF-P4 · ACF-P5 · ACF-P6 · ACF-P7 · ACF-P8 · ACF-P9 · ACF-P10 · ACF-P11 · ACF-P12 · ACF-P13 · ACF-P14 · ACF-P15


Read the worked example before the blank form. Every card comes three times: blank, then filled in by Lendari, then by Rhena. The two examples answer the same questions in two industries and at two scales, and read side by side. A card whose two answers you have read fills in twice as fast.

Autonomy **without governance**

An AI agent no longer merely suggests: it **decides and acts on its own**, it approves a loan, screens an application, arbitrates a price, commits a spend on behalf of an organization. **This is not a chatbot**: it is an actor that creates binding commitments.



DAMAGE IN SECONDS · across four fronts at once



Accountability cannot be delegated to a machine



The agent acts · but it is the **company and its executives** who answer for it. The **AI Act** sanctions the company (up to **7% of global annual turnover**); the executive's **personal** liability, both civil and criminal, falls in addition under **national law**.

AI Act

REG. (EU) 2024/1689



The challenge: framing autonomy **without killing it**. Neither **banning** it · value is lost. Nor **leaving it unrestrained** · control is lost. The right answer is not *less* autonomy, it is **governed autonomy**.

Who decides **when the agent decides?**

This is the question **ACF** answers · on the next page.

What is ACF?

The **Agentic Commerce Framework**® is the first governance framework dedicated to **agentic commerce** : the moment when AI agents stop suggesting and start **acting** on behalf of an organization: approving a loan, screening an application, arbitrating a price, committing a spend. This “commerce” does not mean e-commerce: it is **any decision delegated to an agent** , from finance to industry, from human resources to public administration.

Who answers when an agent gets it wrong?

The law exists · the **AI Act** mandates human oversight · but no **operational playbook** linked the machine's decision to the accountability of a named human. ACF fills that gap.

THE DEFINITION

A framework that lets you **delegate decisions to AI agents** without losing control, compliance, or accountability.

THE 4 FOUNDING PRINCIPLES

FOUNDATION

Decisional sovereignty

You retain **command** over your decisions, your data, your dependencies.

ACF-01

DESIGN

Governance by design

Rules **before** deployment: mandates, thresholds, scopes.

ACF-02 · 03 · 04 · 13

CONTROL

Ultimate human control

The helm is never relinquished: **oversee, take over, stop**. AI Act Art. 14.

ACF-06 · 07 · 12

EVIDENCE

Traceable accountability

Every decision remains **attributable to a named human**, the DDAO (Delegated Decision Agent Officer).

ACF-09 · 17

No **autonomy** without supervision · no supervision without governance · no governance without **sovereignty**.

A third way · neither deregulation nor state control: **sovereignty**.

For **executives, compliance teams**, the **DDAO, educators** and their **students**.

The European third way

AI is not just another product: it is a **metatechnology** that redistributes the power to decide. Two models dominate the response : the platform and the State. Europe is sketching a third way: **sovereignty through law and through design**. ACF is its concrete implementation.

In the tradition of the work of **ISN** (Institute for Digital Sovereignty) × **iDFRights** · Benhamou, Cavada · 2026

*"AI is not just a tool or a platform, it is a **transformative metatechnology** : the technology behind technology and everything else, itself a producer of tools and platforms."*

Mustafa Suleyman · The Coming Wave, 2023



PLATFORM MODEL

Innovation through concentration

Speed, capital and network effects. Powerful agents, deployed fast, at global scale.

RISK dependency, capture of decision-making, eroded sovereignty.



CENTRALIZED MODEL

Control through planning

AI steered by public authority, vertically integrated, placed at the service of order.

RISK surveillance, rigidity, subordination of the individual to the system.



ACF

THIRD WAY

Sovereignty through law and design

GDPR, AI Act, **decision sovereignty**. Autonomy is delegated, but remains governed and reversible.

REQUIREMENT equip the doctrine with tools · otherwise it remains a wish.



The third way has its doctrine · **ISN, iDFRights, AI Act** , but it lacks the field. ACF delivers what turns a principle into practice: **an open standard, tools, cases, a simulation**. The framework is free for schools and training programs; the enterprise ecosystem remains under contract. **Open standard, protected value**.

Every domain has its **reference framework**

Every major digital challenge has eventually acquired its own standard. Security has ISO 27001, personal data has the GDPR, AI management has ISO 42001, high-risk AI has the AI Act. One domain remained without a dedicated framework: **the governance of decisions delegated to agents**.

Cybersecurity**ISO 27001****Data protection****GDPR****AI management****ISO 42001****High-risk AI****AI Act****Agentic governance** ACF'S DOMAIN**ACF**

ACF is not just one more add-on tool. It is the **reference framework** for a domain that did not yet officially exist: **agentic governance**. Naming this domain and giving it its discipline is already to begin instituting it.

The 4 new laws of robotics

Where Asimov wrote to protect humans from machines, **Frank Pasquale** writes to protect human societies from the **concentration of power** that AI makes possible. ACF does not merely cite them: it **wires** them into enforceable tools.

Frank Pasquale · **New Laws of Robotics** · Harvard University Press, 2020

LAW

1

AI must complement professionals, never replace them.

IN ACF The agent **proposes**, the human retains the decision. Autonomy levels (ACF-02) bound delegation decision by decision; arbitration and correction remain a human act.

LAW

2

AI must not counterfeit humanity.

IN ACF Every agent declares itself as an agent. The **Agent Card (ACF-05)** sets its identity and scope; no system may impersonate a human to a customer or a colleague.

LAW

3

AI must not intensify the arms race.

IN ACF Criticality (ACF-03) and risk assessment (ACF-12) bound escalation. Negative-sum uses are refused: delegated power serves a sustained objective, not one-upmanship.

LAW

4

AI must always disclose its creators, controllers, and owners.

IN ACF Each agent has a named human owner and an accountable **DDAO** (Delegated Decision Agent Officer). The Decision Registry (ACF-09) traces the origin of every action : traceability is designed in from the start, not bolted on after the fact.



In ACF, these four laws are not an ethical add-on. They are inscribed in the **Agentic Constitution (ACF-04)** as non-negotiable principles, and made enforceable through **Accountability by Design (ACF-14)**: a principle that does not translate into a tool binds no one.

The 4 founding principles

The entire framework derives from four founding principles. These are not displayed values: they are the **arbitration rules** that decide when to delegate, how far, and under which conditions. Every card in the toolkit is a concrete implementation of them.

1



Decision-making sovereignty

Critical strategic decisions are **never delegated** to autonomous agents. Autonomy can be delegated; sovereignty cannot be ceded.

2



Governance by design

The governance framework is defined **before deployment**, not after the incident. Mandates, thresholds, and safeguards precede production rollout.

3



Ultimate human control

Any agentic system **preserves a capacity for human intervention** : supervision, takeover, and up to the kill switch.

4



Traceable accountability

Every autonomous action is **auditable and attributable** to an identified owner. Who decided, on what basis, at what moment.



These four principles are the **"why"** of the framework. The **five pillars of Accountability by Design** : Identification, Logging, Recourse, Anticipation, Reversibility (card ACF-14) · are its operational **"how"**.

The 4 layers of agentic governance

ACF structures the delegation to agents in four nested layers: from strategic sovereignty down to supervised execution. Every decision flows down the stack, and never reaches the real world without passing through an arbitration point.

LAYER

1



Governance & Sovereignty

Who decides, who answers, in the name of which values. Mandates, roles (including the DDAO, Delegated Decision Agent Officer), sovereignty of data and choices. The layer that makes everything else enforceable.



LAYER

2



Decision policy

The written rules: what can be delegated to agents, under which criticality threshold, at which autonomy level. Doctrine translated into enforceable policy.



LAYER

3



Agents

The autonomous agents that perceive, reason and **propose** decisions, within the perimeter set by the upper layers.



ENTERPRISE BUILDING BLOCK · ACF ECOSYSTEM

ACF Decision Engine



Arbitration point between the agent and the real world: **no sensitive decision is executed without validation**. Available under contract for organizations · inner workings not disclosed publicly.



ACFSTANDARD.COM



LAYER

4




Execution & supervision

Action in the real world, traced and monitored: logging, alerts, kill switch. The layer that executes · and that is held accountable.

The 4 levels of delegated autonomy

Delegation is not all or nothing. For **each type of decision**, ACF defines an autonomy level : from a simple suggestion to autonomous action. The higher you go, the more the agent decides on its own; the tighter upstream governance must be.



HUMAN CONTROL


AGENT AUTONOMY




The authorized level is not a matter of comfort: it **derives from the criticality score**. The more critical a decision (**irreversible, high-impact**), the lower the permitted autonomy level : until it reverts entirely to the human beyond the threshold.

Measuring the **criticality** of a decision

Before delegating a decision to an agent, ACF measures it across three axes · **impact**, **frequency**, **irreversibility** : each scored from 1 to 10. The product, divided by 10, yields a **criticality score** that determines the governance regime.



THE SCORE POSITIONS THE DECISION ON THE ACF SCALE




The threshold is not an "extra" risk rating: it is a **governance routing rule**. Two decisions with identical impact may end up in two different regimes if one is **reversible** and the other is not. Irreversibility is what drives the score up, and what brings the human back into the loop.

The risk matrix

Every operational risk is plotted along two axes: its **likelihood** and its **severity**. Their product yields a score from 1 to 25 that places the risk within one of four bands · and sets the urgency of remediation.

Score = Likelihood × Severity.



Do not confuse the two: the risk matrix (ACF-12) assesses an **operational risk** to be remediated; the **criticality score** (ACF-03) determines the **level of autonomy permitted** to an agent. Two distinct scales, two distinct uses · often cross-referenced, never merged.

The 4 levels of agentic maturity

Where the 4 autonomy levels qualify **a single decision**, maturity qualifies **an agent as a whole**: how far has the organization handed over the wheel? ACF distinguishes four tiers, from mute automation to supervised autonomy.



HUMAN CONTROL

AGENT AUTONOMY



Do not conflate the two scales: **maturity** (N0→N3) describes where **an agent stands overall**; the **4 autonomy levels** (Suggestive → Autonomous) describe what it is allowed to do **decision by decision**. An N2 agent may remain "Suggestive" on critical decisions.

The DDAO guardian of delegation

Delegating to an agent never dilutes accountability: it concentrates on a named individual, the **DDAO** (Delegated Decision Agent Officer). This is the human who answers for delegated decisions : the one who sets the frame, monitors, and retains the power to halt everything.

D

Delegated

The decision is delegated, never abandoned.

D

Decision

Their scope: the decisions entrusted to the agent.

A

Agent

The agentic agent operating under mandate.

O

Officer

An identified, accountable human owner.

THEIR MISSIONS



Define the mandates

Scope, thresholds and safeguards for each agent. [ACF-13](#)



Set the autonomy levels

What the agent may decide alone, according to criticality. [ACF-02 · 03](#)



Supervise continuously

Monitor decisions, retain the ability to take back control. [ACF-06](#)



Hold the kill switch

Power to suspend the agent immediately, without condition. [ACF-07](#)



Approve audits and plans

Sign off on governance audits and corrective actions. [ACF-10 · 11](#)



Answer for decisions

Report to leadership: who, on what basis, when. [ACF-14](#)



The DDAO embodies the first pillar of **Responsibility by Design**: **identification**. An autonomous action without a named owner is not delegation · it is an abdication of sovereignty.

Deploying ACF in 6 steps

ACF is not a one-time checklist. It is a **cycle**: diagnose, map, frame, deploy, supervise, audit, then start over. Here is the order of operations, and the toolkit cards that equip each step.

01

Diagnose

Measure your exposure and digital dependency before any delegation. The mandatory starting point.

ACF-01

ACF-02

ACF-05

ACF-06

ACF-07

ACF-P11

02

Map

List the decisions, score their criticality, and identify those that can never be delegated.

ACF-02

ACF-03

ACF-01

ACF-03

03

Frame

Write the rules: Agentic Constitution, agent scope, mandate, thresholds, and safeguards.

ACF-04

ACF-05

ACF-13

ACF-15

ACF-04

ACF-P1

ACF-P5

ACF-P9

ACF-P10

ACF-P12

04

Deploy

Roll out a first governed agent within a controlled and reversible perimeter.

ACF-08

ACF-P2

ACF-P3

05

Supervise

Monitor continuously, log decisions, stay in control · all the way to the kill switch.

ACF-06

ACF-07

ACF-09

ACF-16

ACF-09

ACF-P4

ACF-P6

ACF-P7

ACF-P8

ACF-P13

ACF-P14

06

Audit

Evaluate, trace accountability, correct. The audit feeds into the next cycle.

ACF-10

ACF-11

ACF-12

ACF-14

ACF-17

ACF-08

ACF-010

ACF-P15



Agentic governance is a **continuous cycle**, not a one-off project. Step 6 feeds Step 1 of the next cycle: every **incident**, **every deviation** refines the mandates, thresholds, and safeguards.



READ FIRST

You are holding a kit of **45 cards** designed to govern AI agents within your organization. This guide explains **where to start, which card to use in which situation, and which path to follow based on your profile.**

1 WHAT IS AN AI AGENT?

An **AI agent** is not an ordinary conversational AI (like ChatGPT, which answers questions). It is a system that **decides** and **acts** autonomously within a defined perimeter.

Conventional generative AI

You ask - the AI **proposes** - you **decide** and execute.
E.g., drafting an email, summarizing a text, brainstorming.

Autonomous AI agent

You grant it a **mandate** - the agent **analyzes, decides and acts** alone within its perimeter.
E.g., screening resumes, adjusting prices, freezing a suspicious transaction.

2 WHY AGENTIC GOVERNANCE

An agent acting autonomously can **drift** - and cause legal, ethical, financial or reputational harm **at machine speed**. Agentic governance frames that autonomy without killing it.

Legal stake

European regulatory framework on AI and data (AI Act, in force since 2024). Heavy administrative penalties for the organization (up to 7% of worldwide turnover), compounded by the personal liability of executives under national law.

Ethical stake

When an agent decides about a person (hiring, credit, healthcare), the question of fairness shifts entirely.

Business stake

Velocity and productivity gained, but also the scale of damage a single incident can cause in seconds, loss of differentiation, dependency on a single vendor.

Democratic stake

Decisional sovereignty of organizations and States.
Who really decides when the agent decides?



3 THE ACF VOCABULARY IN 18 CONCEPTS

The eighteen notions to master before working with the toolkit. Each one is anchored in one or more cards.

DECISIONAL SOVEREIGNTY

ACF-01

An organization's capacity to know, at any moment, who actually decides - when agents are acting.

AI AGENT ACF-05

A system that decides and acts autonomously within a defined perimeter. Not a mere chatbot.

DELEGATED AUTONOMY ACF-02

4 levels set **decision by decision**:
Advisory · Co-decision · Supervised autonomous · Autonomous.

MANDATE ACF-13

The document defining an agent's perimeter, limits and accountabilities - comparable to a job description.

AGENTIC CONSTITUTION ACF-04

Non-negotiable foundational rules governing an agent's behavior. Signed by 4 accountable officers.

GUARDRAILS ACF-07

Mechanisms that prevent drift: validation thresholds, escalations, emergency stop (Kill switch).

DECISION REGISTER ACF-09

A log one can **only add to** (never erase or modify). All agentic decisions are recorded - auditable and legally enforceable.

DDAO ACF-14

Delegated **D**ecision **A**gent **O**fficer. The role that carries responsibility-by-design, with a reasoned right of objection.

RESPONSIBILITY BY DESIGN

ACF-14

Pasquale principle (Harvard, 2020), *responsibility by design*: accountability is built into the system from its inception, not bolted on afterward.

CRITICALITY ACF-03

What determines whether a decision may be delegated. It is calculated, not estimated: impact, reversibility, reach.

AGENTIC MATURITY ACF-05

Four levels, N0 to N3, set **once for the agent as a whole**. Not to be confused with delegated autonomy, which is set decision by decision.

SHADOW AGENT ACF-06

An agent running in the organization that governance never recorded. It still decides, and nobody answers for it.

AGENT FLEET ACF-09

Every agent in service. Past the first one, governance no longer holds by attention; it holds by inventory.

NON-DELEGABLE ZONE ACF-08

What an agent will never decide, whatever its level. Shutting the agent down is part of it: it does not decide its own extinction.

ESCALATION ACF-P6

The written path by which the agent hands back to a named human, with the deadline and the reason. Without it, a block becomes an incident.

COUNTER-METRIC ACF-15

The figure that contradicts the objective. An agent that hits its target by degrading something else has not succeeded: it has moved the problem.

DEGRADED MODE ACF-P12

How the work carries on when the agent stops. Without a degraded mode, a kill switch is a button nobody will dare press.

HOLDING THE RECORD ACF-P14

Who keeps the logs, where, and for how long. A record held by the vendor alone is not evidence; it is a dependency.



4 THE 45 CARDS, BY FAMILY AND BY USE

The 45 cards fall into **3 families**: Decide (ACF-01 to ACF-20), Organise (ACF-O1 to ACF-O10) and Operate (ACF-P1 to ACF-P15). Decide establishes what is delegated and within which limits, Organise looks at what the agent shifts inside the organization, Operate deals with what happens once it is live. A card may serve several uses; these are indicative.

DIAGNOSE

ACF-01 Sovereignty score

FRAMING

ACF-02 Decision map · **ACF-03** Criticality matrix ·
ACF-04 Agentic constitution · **ACF-05** Agent card ·
ACF-13 Agent mandate · **ACF-15** Objective and
counter-metric

DEPLOYMENT

ACF-08 First agent briefing

RUNNING

ACF-06 Supervision & governance · **ACF-07** Kill switch
· **ACF-09** Decision register · **ACF-16** Supervision test

AUDIT & IMPROVEMENT

ACF-10 Action plan · **ACF-11** 30-day governance audit
· **ACF-12** Agentic risk assessment · **ACF-14**
Responsibility by design · **ACF-17** Validity review

TEACHING

ACF-18 Guided case study · **ACF-19** Instructor's guide
· **ACF-20** Governance simulation

THE HEART OF THE TOOLKIT

If you were to keep only **3 cards**: ACF-01 (diagnose), ACF-13 (frame an agent by mandate), ACF-09 (record the decisions). Together the three cover 80% of basic governance needs.



4 ORGANISE: WHAT THE AGENT SHIFTS INSIDE THE COMPANY

Decide says what is entrusted to the agent. **Organise** (ACF-O1 to ACF-O10) looks elsewhere: who loses control, which skill stops being maintained, which agents run without anyone having declared them. These questions always come up after the first deployment.

WHO DOES WHAT, FROM NOW ON

ACF-O1 Human-AI-agent matrix · **ACF-O2** Hybrid organization map · **ACF-O3** Critical skills to keep

THE RULES OF THE HOUSE

ACF-O4 AI policy · **ACF-O8** Exceptions register

WHAT ESCAPES CONTROL

ACF-O5 Shadow AI · **ACF-O6** Shadow agents · **ACF-O7** The shift of power

BEYOND THE FIRST AGENT

ACF-O9 Fleet governance · **ACF-O10** Affected people and redress

WHERE TO ENTER

Start with ACF-O6: you cannot govern a fleet you have not inventoried. Then ACF-O4, which writes the rule, and ACF-O8, which records every departure from it.



4 OPERATE: WHAT HAPPENS ONCE THE AGENT IS LIVE

Operate (ACF-P1 to ACF-P15) takes over the day the agent runs for real. Configuration, acceptance testing, the weekly supervision meeting, the incident, the vendor outage: everything the deployment dossier could not anticipate.

GOING LIVE

ACF-P1 Configuration grid · **ACF-P2** Acceptance protocol · **ACF-P3** Ramp-up plan

DAY-TO-DAY RUNNING

ACF-P4 Supervision log · **ACF-P5** Role allocation · **ACF-P6** Escalation & trust

WHEN IT BREAKS

ACF-P7 Stop test · **ACF-P8** Incident log · **ACF-P12** Degraded mode & exit

DEPENDENCIES AND EVIDENCE

ACF-P9 Permissions review · **ACF-P10** Vendor assessment · **ACF-P11** Foundation dependency · **ACF-P13** Model version log · **ACF-P14** Holding the record · **ACF-P15** Blind spot statement

WHERE TO ENTER

ACF-P4 holds supervision over time, ACF-P12 says how the work carries on when the agent stops, and ACF-P14 guarantees evidence still exists the day it is asked for.



5 @ CHOOSE YOUR LEARNING PATH

Five paths tailored to your profile. Each one proposes an **ordered sequence of cards** to work through - not a mandate, simply a recommendation.

A Executive discovering agentic governance

90 MIN

You want to grasp the issues quickly to decide whether your organization should equip itself.

ACF-01 Score →

ACF-02 Decision map →

ACF-03 Criticality →

ACF-07 Shift of power →

ACF-11 30-day audit →

ACF-14 DDAO

B Professor or trainer preparing a course

3-4 H

You are about to deliver a course, workshop or executive education program on AI governance.

ACF-19 Instructor's guide →

ACF-01 Score →

ACF-03 Criticality →

ACF-04 Constitution →

ACF-07 Kill switch →

ACF-06 Shadow agents →

ACF-18 Case study →

ACF-20 Simulation



5 @ CHOOSE YOUR LEARNING PATH

C Employee tasked with deploying or supervising an AI agent

2-3 H

Your leadership has entrusted you with an AI-agent project: you must frame it properly, then hold it once it runs.

[ACF-08 First-agent dossier →](#)

[ACF-05 Agent card →](#)

[ACF-13 Mandate →](#)

[ACF-P1 Configuration →](#)

[ACF-P2 Acceptance →](#)

[ACF-P3 Ramp-up →](#)

[ACF-07 Kill switch →](#)

[ACF-P4 Supervision →](#)

[ACF-09 Register](#)

D Organization already running several agents

1 DAY

The first agent went live long ago, others followed, and nobody quite knows which ones are running or who answers for them.

[ACF-06 Shadow agents →](#)

[ACF-09 Fleet →](#)

[ACF-04 AI policy →](#)

[ACF-P9 Permissions →](#)

[ACF-P12 Degraded mode →](#)

[ACF-P14 Holding the record](#)

E Independent student or self-learner

2 WEEKS

You are exploring the subject on your own, for a thesis, a project or a certification, and you want to cover all of it.

[ACF-01 Score →](#)

[ACF-02 to 20 Decide →](#)

[ACF-O1 to O10 Organise →](#)

[ACF-P1 to P15 Operate →](#)

[ACF-18 Case study](#)



6 Q WHICH CARD FOR WHICH NEED

DECIDE · ACF-01 TO ACF-20

NEED OR SITUATION	CARD	TYPICAL DURATION
Assess my dependency on external tools and vendors	ACF-01	45 min - 1 h 30
Map the decisions to be delegated across my organization	ACF-02	2 h - 4 h
Decide which level of autonomy to grant to a decision	ACF-03	1 h 45 - 3 h 30
Write an agent's non-negotiable rules	ACF-04	3 h 30 - 6 h 30
Document a specific agent, technically	ACF-05	3 h - 6 h
Set up continuous supervision of an agent	ACF-06	3 h - 5 h 30
Design an emergency stop mechanism	ACF-07	3 h - 6 h
Document the deployment of my first agent	ACF-08	1 - 2 days
Record agentic decisions in an auditable way	ACF-09	1 h 45 - 3 h
Plan corrective actions after an audit or an incident	ACF-10	3 h 30 - 6 h
Audit governance maturity in 30 days	ACF-11	1 - 1.5 days
Assess the risks an agent creates	ACF-12	1 - 2 days
Issue a formal mandate to an agent	ACF-13	1 - 1.5 days
Audit a system's responsibility by design	ACF-14	2 h - 4 h
Set an agent's objective and the figure that contradicts it	ACF-15	3 h 30 - 6 h 30
Check that supervision actually works	ACF-16	3 h - 5 h
Reopen a delegation that has lapsed	ACF-17	3 h 30 - 6 h
Build a teaching case study	ACF-18	0.5 - 1.5 days
Run a workshop or a course on ACF	ACF-19	3 h
Run a governance simulation	ACF-20	4 h



6 🔍 WHICH CARD FOR WHICH NEED (CONT.)

ORGANISE · ACF-O1 TO ACF-O10

NEED OR SITUATION	CARD	TYPICAL DURATION
Split the work between humans, AI and agents	ACF-O1	1 h 15 - 2 h 30
Map the organization once the agents are in place	ACF-O2	1 h 45 - 3 h
Name the skills that must not be allowed to lapse	ACF-O3	1 h - 2 h
Write the company's AI policy	ACF-O4	1 h 15 - 2 h 30
Inventory the AI uses nobody declared	ACF-O5	1 h - 2 h
Inventory the agents nobody declared	ACF-O6	1 h 15 - 2 h 30
See where decision power has shifted	ACF-O7	45 min - 1 h 45
Keep the register of departures from the rule	ACF-O8	45 min - 1 h 30
Govern a fleet of agents rather than a single one	ACF-O9	1 h 15 - 2 h 30
Open a redress route to the people a decision affects	ACF-O10	45 min - 1 h 45



6 Q WHICH CARD FOR WHICH NEED (CONT.)

OPERATE · ACF-P1 TO ACF-P15

NEED OR SITUATION	CARD	TYPICAL DURATION
Configure an agent, line by line	ACF-P1	2 h - 4 h
Acceptance-test an agent before it goes live	ACF-P2	45 min - 1 h 30
Open access gradually rather than all at once	ACF-P3	1 h - 1 h 45
Hold the weekly supervision meeting	ACF-P4	45 min - 1 h 30
State who does what around the agent	ACF-P5	1 h 30 - 2 h 30
Write how the agent hands back to a human	ACF-P6	30 min - 1 h 15
Test the stop mechanism for real	ACF-P7	1 h 30 - 3 h
Learn from an incident without spending the week on it	ACF-P8	1 h 30 - 2 h 30
Review what the agent can actually reach	ACF-P9	1 h - 1 h 45
Assess the vendor behind the agent	ACF-P10	1 h - 2 h
Map what the agent depends on to function	ACF-P11	1 h 30 - 3 h
Plan how the work carries on without the agent	ACF-P12	2 h - 4 h
Track model version changes	ACF-P13	1 h - 2 h 30
Decide who keeps the records, where, and for how long	ACF-P14	1 h 45 - 3 h 30
Declare what the agent cannot see	ACF-P15	1 h 45 - 3 h



7 >> TYPICAL DEPLOYMENT WORKFLOW

The logical sequence for deploying an agent within an organization. The cards build on each other - every stage is documented.



AND AFTER GOING LIVE

Deployment is not the end of the work; it is the moment the work changes nature. This second sequence is the day-to-day life of an agent, then of a fleet: it draws on the Organise and Operate families.





8 ? FREQUENTLY ASKED QUESTIONS

Do I need to complete all 45 cards to use it properly?

No. The toolkit is a **menu**, not a checklist. Depending on your situation, 3 to 8 cards will suffice. The table “which card for which need” tells you which one.

May I adapt the cards to my context?

Yes - provided ACF is credited and no commercial use is made without agreement. See the usage charter for the full scope of permissions.

Who is the DDAO (Delegated Decision Agent Officer) role designed for?

A senior profile, age 35-65, with dual experience (business + tech, or business + law). Profiles observed in practice: former regulator, former CISO, former AI counsel, former AI Safety researcher. Capable of saying NO to executive leadership and defending that position before the Board.

Is the toolkit alone sufficient to achieve regulatory compliance?

No - it is an operational and pedagogical framework, not a substitute for legal counsel. The cards help structure the approach, but legal expertise remains indispensable for systems classified as high-risk under European regulation.

How do I contribute or suggest improvements?

Via the website acfstandard.com (contact area). Subsequent editions of the toolkit integrate field feedback - with attribution if you wish.

TO GO FURTHER

The website · acfstandard.com - instructor area, updates, supplementary resources.

ACF in Action · 18 analyses of agentic decisions actually taken by named organizations - acfstandard.com/en/acf-action.

ACF Playbook · sector handbooks running the full path on fictional but recognizable organizations, forthcoming.

Teaching cases · 15 cases, 3 per program track (Business, Engineering, MBA, EMBA, Master in AI) - included with the instructor pack.

Slide deck · 28 introductory slides on AI-agent governance and the European regulatory framework, ready to project.

Grading rubric · Auto-calculated Excel sheet for evaluating student deliverables on a 100-point scale.

Decide

The first family, and the only one no organisation can skip. It establishes what is delegated, within which limits, under whose authority, and by which mechanism control is taken back. Without it there is nothing to organise and nothing to operate.

THE QUESTION THIS FAMILY ASKS

Who authorised this decision, and how far does it go?

THE 20 CARDS

ACF-01	Sovereignty Score	ACF-08	First Agent Briefing	ACF-15	Objective and Counter-Metric
ACF-02	Decision Map	ACF-09	Agentic Decisions Register	ACF-16	Supervision Test
ACF-03	Criticality Matrix	ACF-10	Action & Improvement Plan	ACF-17	Validity Review
ACF-04	Agentic Constitution	ACF-11	30-Day Governance Audit	ACF-18	Guided Case Study
ACF-05	Agent Card	ACF-12	Agentic Risk Assessment	ACF-19	Instructor's Guide
ACF-06	Supervision & Governance	ACF-13	Agent Mandate	ACF-20	Governance Simulation
ACF-07	Kill Switch	ACF-14	Responsibility by Design		

WHAT HAPPENS WHEN THIS FAMILY IS MISSING

- Nobody knows who authorised the agent to decide, and the question only comes up on the day of the incident (**ACF-13**).
- The limits live in people's heads rather than in a document anyone can produce (**ACF-04**).
- The stop mechanism has never been tried, and you find out it does not work at the moment you need it (**ACF-07**).

WHERE TO OPEN IT

ACF-01	What you already delegate, without having decided to.
ACF-02	Which decisions run through the agent, and at what level of autonomy.
ACF-13	The mandate, written, signed, enforceable.

The last three cards are not working tools but the documents of a session: a guided case for the student, a guide for the instructor, a simulation for the group.



OBJECTIVE

Before agentifying anything, measure your degree of **dependence**, and therefore your exposure : across four dimensions. The greater the dependence, the weaker the sovereignty: **govern before you automate**.

1 ⌚ HOW TO READ THIS SCORE

Rate each dimension from **0 to 25** based on your level of dependence (0 = fully autonomous, 25 = fully dependent). The sum yields a **dependence index out of 100**: this is the inverse of your sovereignty.

Scoring guide : for each dimension, count the number of questions you answer « yes » to:

0 – 8 none · you retain control **9 – 16** one · partial dependence **17 – 25** both · strong dependence

2 ≡ THE FOUR DIMENSIONS OF DEPENDENCE

TECHNOLOGICAL

/ 25

Actual control over your models, infrastructure, and data.

- Would a change of provider (pricing, terms of service, shutdown) paralyze you?
- Does your sensitive data flow outside your control?

YOUR SCORE / 25

ACQUISITION

/ 25

Share of your customers reached through channels you do not own (e.g. marketplaces).

- Does a third-party platform capture the bulk of your customers?
- Would delisting cut you off from the market?

YOUR SCORE / 25

DECISIONAL

/ 25

Decisions already delegated to systems without human governance.

- Are binding decisions being made without review?
- Could you justify a decision made by your AI?

YOUR SCORE / 25

FINANCIAL

/ 25

Concentration of your revenue or funding among a small number of players.

- Does a single customer or channel account for a critical share of revenue?
- Would its loss threaten business continuity?

YOUR SCORE / 25



3 YOUR DEPENDENCE INDEX

Add up your four scores. **The higher the index, the greater your exposure** , and the more governance must precede automation. / 100

0 – 33

SOVEREIGNTY UNDER CONTROL

Agentification is possible with light-touch governance. Proceed to ACF-02.

34 – 66

MODERATE EXPOSURE

Govern and reduce dependencies before expanding agent autonomy.

67 – 100

CRITICAL DEPENDENCE

Reduce your exposure **before** any further automation. Top priority.

4 YOUR THREE ACTION PRIORITIES

Identify the dimension where you are most dependent: that is where your sovereignty is at stake. Note three concrete actions to reduce it.

#	ACTION TO REDUCE THE DEPENDENCE	TARGET DIMENSION	DEADLINE
1			
2			
3			

Next step: once your exposure is measured, map your key decisions with **ACF-02 · Decision Map**, then assess their criticality with **ACF-03**.



RUNNING EXAMPLE

Lendari : online fashion brand (€120M revenue, 90 employees, Paris). Project: deploying "**Margaux**", a customer service & returns agent. Diagnostic led by **Camille Roussel**, Chief Digital Transformation Officer (DDAO · Delegated Decision Agent Officer).

PAGE 1/2

1 HOW TO READ THIS SCORE

Rate each dimension from **0 to 25** based on your level of dependency (0 = fully autonomous, 25 = fully dependent). The sum gives you a **dependency index out of 100**: the inverse of your sovereignty.

Rating guide : for each dimension, count the number of questions you answer "yes" to:

0 – 8 none · you retain control **9 – 16** one · partial dependency **17 – 25** both · strong dependency

2 THE FOUR DIMENSIONS OF DEPENDENCY

TECHNOLOGICAL

/ 25

Actual control over your models, infrastructure, and data.

- Would a change of vendor (pricing, terms, shutdown) paralyze you?
- Does your sensitive data transit outside your control?

LENDARI SCORE **15** / 25

ACQUISITION

/ 25

Share of your customers acquired through channels you do not own (e.g. marketplaces).

- Does a third-party platform capture the bulk of your customers?
- Would delisting cut you off from the market?

LENDARI SCORE **20** / 25

DECISIONAL

/ 25

Decisions already delegated to systems operating without human governance.

- Are binding decisions made without validation?
- Could you justify a decision made by your AI?

LENDARI SCORE **13** / 25

FINANCIAL

/ 25

Concentration of your revenue or financing on a small number of players.

- Does a single customer or channel account for a critical share of revenue?
- Would its loss threaten business continuity?

LENDARI SCORE **12** / 25



3 YOUR DEPENDENCY INDEX

Add up your four scores. **The higher the index, the greater your exposure** , and the more governance must precede automation.

15 + 20 + 13 + 12 **60 / 100**

0 – 33

SOVEREIGNTY UNDER CONTROL

Agentification possible with light governance. Move on to ACF-02.

34 – 66

MODERATE EXPOSURE

Govern and reduce dependencies before expanding agent autonomy.

◀ LENDARI: 60 / 100

67 – 100

CRITICAL DEPENDENCY

Reduce exposure **before** any further automation. Absolute priority.

4 YOUR THREE ACTION PRIORITIES

Identify your most dependent dimension: that is where your sovereignty is at stake. For Lendari, it is **Acquisition** (20/25) · hence its priorities.

#	ACTION TO REDUCE DEPENDENCY	TARGET DIMENSION	DEADLINE
1	Build owned audience (CRM, mobile app, SEO, loyalty) to reduce the share of paid channels	Acquisition	9 months
2	Regain control of customer & returns data on a controlled EU instance	Technological	4 months
3	Govern Margaux's decisions (constitution, supervision, kill switch) before any increase in autonomy	Decisional	Immediate

Next step: once your exposure is measured, map your key decisions with **ACF-02 · Decision Map**, then assess their criticality with **ACF-03**.



RUNNING EXAMPLE

Rhena : German insurer (€75bn in premiums, 62,000 employees, 28 countries, Cologne). Project: deploying "**ANNA**", a level 1 home claims agent. Diagnostic led by **Dr Katrin Ostermann**, Chief Compliance Officer (DDAO · Delegated Decision Agent Officer).

PAGE 1/2

1 HOW TO READ THIS SCORE

Rate each dimension from **0 to 25** based on your level of dependency (0 = fully autonomous, 25 = fully dependent). The sum gives you a **dependency index out of 100**: the inverse of your sovereignty.

Rating guide : for each dimension, count the number of questions you answer "yes" to:

0 – 8 none · you retain control **9 – 16** one · partial dependency **17 – 25** both · strong dependency

2 THE FOUR DIMENSIONS OF DEPENDENCY

TECHNOLOGICAL

/ 25

Actual control over your models, infrastructure, and data.

- Would a change of vendor (pricing, terms, shutdown) paralyze you?
- Does your sensitive data transit outside your control?

RHENA SCORE **16** / 25

ACQUISITION

/ 25

Share of your customers acquired through channels you do not own (e.g. marketplaces).

- Does a third-party platform capture the bulk of your customers?
- Would delisting cut you off from the market?

RHENA SCORE **6** / 25

DECISIONAL

/ 25

Decisions already delegated to systems operating without human governance.

- Are binding decisions made without validation?
- Could you justify a decision made by your AI?

RHENA SCORE **17** / 25

FINANCIAL

/ 25

Concentration of your revenue or financing on a small number of players.

- Does a single customer or channel account for a critical share of revenue?
- Would its loss threaten business continuity?

RHENA SCORE **9** / 25



3 YOUR DEPENDENCY INDEX

Add up your four scores. **The higher the index, the greater your exposure** , and the more governance must precede automation.

48 / 100

16 + 6 + 17 + 9

0 – 33

SOVEREIGNTY UNDER CONTROL

Agentification possible with light governance. Move on to ACF-02.

34 – 66

MODERATE EXPOSURE

Govern and reduce dependencies before expanding agent autonomy.

◀ RHENA: 48 / 100

67 – 100

CRITICAL DEPENDENCY

Reduce exposure **before** any further automation. Absolute priority.

4 YOUR THREE ACTION PRIORITIES

Identify your most dependent dimension: that is where your sovereignty is at stake. For Rhena, it is the **Decision** dimension (17/25) · hence its priorities.

#	ACTION TO REDUCE DEPENDENCY	TARGET DIMENSION	DEADLINE
1	Log the depreciation schedule version on every decision and tie it to the tariff revision	Decisional	Immediate
2	Extend stop authority in writing to the US subsidiary and test it on both platforms	Decisional	2 months
3	Qualify a second vendor and contract an exit notice period	Technological	9 months

Next step: once your exposure is measured, map your key decisions with **ACF-02 · Decision Map**, then assess their criticality with **ACF-03**.



OBJECTIVE

Map your key decisions and set, for each one, the **level of autonomy** you delegate to the agent. This is the bridge between the sovereignty diagnostic (**ACF-01**) and the criticality analysis (**ACF-03**).

1 THE AGENTIC MATURITY SCALE

Before mapping, establish a shared vocabulary. First, position the maturity targeted for your agent : most organizations aim for level 2. Delegated autonomy is then set decision by decision (section 3).

0

Classical automation

Fixed rules, no autonomy of judgment. A script, not an agent.

1

Assisted agents

The agent proposes, the human systematically decides. Decision support.

2

Governed agents

The agent acts within a defined framework, under active supervision.

TARGET FOR ~80% OF CASES

3

Supervised autonomy

The agent decides alone; the human checks after the fact and keeps the kill switch.

2 YOUR KEY DECISIONS

Identify the **three decisions** you plan to entrust : fully or in part : to an agent. Describe them before assessing their criticality.

A DECISION A

TITLE

MAIN STAKE

EXPECTED IMPACT

KEY DEADLINE

B DECISION B

TITLE

MAIN STAKE

EXPECTED IMPACT

KEY DEADLINE

C DECISION C

TITLE

MAIN STAKE

EXPECTED IMPACT

KEY DEADLINE



3

≡

DETAILED MAPPING

For each decision, record its criticality (calculated in **ACF-03**), the delegated autonomy level (suggestive · co-decision · supervised autonomous · autonomous) and the safeguards that make this level acceptable.

DECISION	CRITICALITY (ACF-03)	DELEGATED AUTONOMY	SAFEGUARDS & CONDITIONS	OWNER
A				
B				
C				

4

✓

THREE ACTIONS TO LAUNCH

#	ACTION TO LAUNCH	RELATED DECISION	DEADLINE
1			
2			
3			

Next step: assess the criticality of each decision with **ACF-03 · Criticality Matrix**, then set the rules of each agent in **ACF-04 · Agentic Constitution**. A well-mapped decision is a decision already half-governed.



RUNNING EXAMPLE

Lendari maps the three decisions entrusted to "Margaux" (customer service & returns agent). **Camille Roussel** (DDAO, Delegated Decision Agent Officer) sets the delegated autonomy level for each.

PAGE 1/2

1 THE AGENTIC MATURITY SCALE

The shared vocabulary of the track. Lendari targets **maturity level 2 (governed agent)** for Margaux: acting within a defined framework, under active supervision. Delegated autonomy is then set decision by decision (section 3).

0 Classical automation Fixed rules, no autonomy of judgment. A script, not an agent.	1 Assisted agents The agent suggests, the human systematically decides. Decision support.	2 Governed agents The agent acts within a defined framework, under active supervision. MARGAUX'S TARGET	3 Supervised autonomy The agent decides alone; the human reviews after the fact and retains the kill switch.
--	---	---	--

2 THE THREE DECISIONS ENTRUSTED TO MARGAUX

A DECISION A	B DECISION B	C DECISION C
TITLE Respond to tier-1 customer service requests (tracking, lead times, availability) KEY ISSUE Relieve the human service desk, answer in < 2 min, 24/7 EXPECTED IMPACT ≈ 60% of tier-1 tickets handled without human intervention KEY MILESTONE Pilot within 2 months	TITLE Authorize a refund or credit note below €100 KEY ISSUE Streamline returns within a controlled financial limit EXPECTED IMPACT Refund time reduced from 5 days to 24h KEY MILESTONE After validation of pilot A	TITLE Offer a goodwill gesture to a dissatisfied customer KEY ISSUE Customer retention · but risk of margin drift EXPECTED IMPACT +5 pts of retention on at-risk customers KEY MILESTONE Phase 2, reinforced supervision



LEVEL

DURATION

USE

TRACK

Foundation

2 h - 4 h

Mapping

Step 2 / 6

RUNNING EXAMPLE

Lendari : detailed mapping of Margaux's decisions and action plan. Criticality scores come from ACF-03.

PAGE 2/2

3 DETAILED MAPPING

For each decision: its criticality (calculated in ACF-03), the delegated autonomy level (suggestive · co-decision · supervised autonomous · autonomous) and the safeguards that make this level acceptable.

DECISION	CRITICALITY (ACF-03)	DELEGATED AUTONOMY	SAFEGUARDS & CONDITIONS	OWNER
A	2.7 ●	supervised autonomous	Sourced answers (customer service knowledge base); escalation as soon as the request falls outside scope.	Margaux · Customer Service Supervisor
B	21 ●	supervised autonomous	Strict cap of €100; systematic logging; daily human review after the fact.	Margaux + customer service validation
C	21 ●	co-decision	Proposal submitted to a human before sending; never any modification of catalog prices.	Customer Service Manager

4 THREE ACTIONS TO LAUNCH

#	ACTION TO LAUNCH	RELATED DECISION	DEADLINE
1	Draft Margaux's constitution (scope, caps, prohibitions) · ACF-04	A · B · C	Immediate
2	Set up daily supervision and the kill switch	B · C	Before the pilot
3	Connect Margaux to an up-to-date customer service knowledge base	A	1 month

Next step: assess the criticality of each decision with ACF-03 · **Criticality Matrix**, then set the rules for each agent in ACF-04 · **Agentic Constitution**. A well-mapped decision is a decision already half governed.



RUNNING EXAMPLE

Rhena maps three home claims decisions examined for "ANNA". Dr Katrin Ostermann (DDAO , Delegated Decision Agent Officer) sets the delegated autonomy level for each.

PAGE 1/2

1 THE AGENTIC MATURITY SCALE

The shared vocabulary of the track. Rhena targets **maturity level 2 (governed agent)** for ANNA: acting within a defined framework, under active supervision. Delegated autonomy is then set decision by decision (section 3).

0 Classical automation Fixed rules, no autonomy of judgment. A script, not an agent.	1 Assisted agents The agent suggests, the human systematically decides. Decision support.	2 Governed agents The agent acts within a defined framework, under active supervision. ANNA'S TARGET	3 Supervised autonomy The agent decides alone; the human reviews after the fact and retains the kill switch.
--	---	--	--

2 THE THREE DECISIONS EXAMINED FOR ANNA

A DECISION A	B DECISION B	C DECISION C
TITLE Acknowledge a claim and chase the missing documents KEY ISSUE Open the file the same day and start the investigation without delay EXPECTED IMPACT ≈ 60% of level 1 files closed with no human rework KEY MILESTONE Pilot on 5% of files from 01/02/2027	TITLE Compute and settle a claim below €2,500 KEY ISSUE Settle on documents, below the mark at which the trade sends an adjuster EXPECTED IMPACT Time to first answer cut from 3 days to under 4h KEY MILESTONE After validation of pilot A	TITLE Decline cover to a policyholder KEY ISSUE A binding decision, reviewable by the Ombudsmann up to €10,000 EXPECTED IMPACT None : the agent builds and reasons the file, it does not rule KEY MILESTONE Outside delegation, whatever the phase



LEVELFoundation

DURATION2 h - 4 h

USEMapping

TRACKStep 2 / 6

RUNNING EXAMPLE

Rhena : detailed mapping of the decisions examined for ANNA and action plan.
Criticality scores come from ACF-03.

PAGE 2/2

3 DETAILED MAPPING

For each decision: its criticality (calculated in ACF-03), the delegated autonomy level (suggestive · co-decision · supervised autonomous · autonomous) and the safeguards that make this level acceptable.

DECISION	CRITICALITY (ACF-03)	DELEGATED AUTONOMY	SAFEGUARDS & CONDITIONS	OWNER
A	3.6 ●	supervised autonomous	Two reminders at most; escalation as soon as the file falls outside level 1.	ANNA · Ines Baumgartner
B	24 ●	supervised autonomous	Strict cap of €2,500; schedule version logged; daily sampled human review.	ANNA + home claims unit review
C	36 ●	suggestive	The agent gathers the documents and cites the clause; it never pronounces the refusal.	Marek Sobczak

4 THREE ACTIONS TO LAUNCH

#	ACTION TO LAUNCH	RELATED DECISION	DEADLINE
1	Draft ANNA's constitution (scope, caps, prohibitions) · ACF-04	A · B · C	Immediate
2	Arm the stop mechanism and test it across both entities	B · C	Before the pilot
3	Tie the depreciation schedule to the yearly tariff revision	A	1 month

Next step: assess the criticality of each decision with ACF-03 · Criticality Matrix, then set the rules for each agent in ACF-04 · Agentic Constitution. A well-mapped decision is a decision already half governed.



OBJECTIVE

Objectively determine, for each type of decision, **how far to delegate it to an agent** : replacing intuition with a reproducible, traceable, and enforceable score, shared across business, technical, and governance teams.

1 THE CRITICALITY FORMULA

The criticality of an agentic decision is calculated along **three axes**, each rated from **1 to 10**. The product is divided by 10 to obtain a readable and comparable score.

$$\text{Score} = (\text{Impact} \times \text{Frequency} \times \text{Irreversibility}) \div 10$$

Each axis is rated from 1 (negligible) to 10 (maximal) · resulting score from 0.1 to 100.

2 THE THREE EVALUATION AXES

IMPACT

1 → 10

Severity of consequences if the decision is wrong. **1** = benign (rephrasing an email). **10** = critical (health, safety, finances, reputation).

FREQUENCY

1 → 10

Recurrence of the decision in operations. **1** = exceptional. **10** = continuous (thousands of times per day).

IRREVERSIBILITY

1 → 10

Difficulty of reversing the decision. **1** = cancellable in a single step. **10** = definitive, no possible return.

3 READING THE SCORE · THE THREE THRESHOLDS



< 15

AGENTIFIABLE

Delegation possible, light supervision. Autonomy set decision by decision (ACF-02).



15 – 25

GOVERNED AGENT

Bounded autonomy: strict policy, traceability, human control.



> 25

HUMAN REQUIRED

The final decision remains human. The agent prepares and proposes · it does not decide.

To combine with: once the zone is determined, set the **delegated autonomy** of each decision (suggestive → autonomous) in the Decision Map **ACF-02**, then record it in the Decision Policy (layer 2) via **ACF-05**.



4 TWO WORKED EXAMPLES

PAGE 2/2

Customer service agent · replying to order tracking questions

$$\begin{array}{|c|} \hline 4 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 9 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 2 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 7.2$$

AGENTIFIABLE

Finance agent · autonomous validation of supplier payments

$$\begin{array}{|c|} \hline 9 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 6 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 8 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 43.2$$

HUMAN REQUIRED

5 YOUR GRID TO COMPLETE

List your candidate decisions, rate each axis from 1 to 10, calculate the score, read the zone. Keep in the governance file as an enforceable record.

DECISION / USE CASE	IMPACT	FREQ.	IRREV.	SCORE	ZONE / VERDICT

6 BEST PRACTICES & POINTS OF VIGILANCE

✓ BEST PRACTICES

- ✓ Rate with multiple stakeholders (business + technical + governance) to neutralize biases
- ✓ Reassess the score with every major evolution of the agent or context
- ✓ Document the "why" of each rating, not just the number
- ✓ When in doubt between two zones, always select the more cautious one

⚠ POINTS OF VIGILANCE

- ! Under-rating irreversibility "because we have a backup" - reputational damage cannot be undone
- ! Confusing low frequency with low risk: a rare case can be catastrophic
- ! Freezing the matrix once and for all · it is a living document
- ! Delegating a red zone "as a pilot" without a real human checkpoint



RUNNING EXAMPLE

Lendari calculates the criticality of "Margaux"'s three decisions. The score determines how far each one can be delegated : the method first, the grid filled in on page 2.

PAGE 1/2

1 THE CRITICALITY FORMULA

The criticality of an agentic decision is calculated along **three axes**, each scored from **1 to 10**. The product is divided by 10 to obtain a readable and comparable score.

$$\text{Score} = (\text{Impact} \times \text{Frequency} \times \text{Irreversibility}) \div 10$$

Each axis is scored from 1 (negligible) to 10 (maximal) · resulting score from 0.1 to 100.

2 THE THREE ASSESSMENT AXES

IMPACT

1 → 10

Severity of consequences if the decision is wrong. **1** = benign (rephrasing an email). **10** = critical (health, safety, finances, reputation).

FREQUENCY

1 → 10

Recurrence of the decision in the activity. **1** = exceptional. **10** = continuous (thousands of times per day).

IRREVERSIBILITY

1 → 10

Difficulty of reversing course. **1** = cancellable in one click. **10** = final, no turning back.

3 READING THE SCORE · THE THREE THRESHOLDS



< 15

AGENTIFIABLE

Delegation possible, light supervision. Autonomy set decision by decision (ACF-02).



15 – 25

GOVERNED AGENT

Bounded autonomy: strict policy, traceability, human oversight.



> 25

HUMAN REQUIRED

The final decision remains human. The agent prepares and proposes · it does not decide.

To be combined with: once the zone is determined, set the **delegated autonomy** of each decision (suggestive → autonomous) in the Decisional Map **ACF-02**, then record it in the Decision Policy (layer 2) via **ACF-05**.



RUNNING EXAMPLE

Lendari : the grid filled in: Margaux's three decisions, three scores. They feed the "Criticality" column of **ACF-02**.

PAGE 2/2

4 THE METHOD, ON TWO EXAMPLES

Customer service agent · answering order tracking inquiries

$$\begin{array}{|c|} \hline 4 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 9 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 2 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 7.2$$

AGENTIFIABLE

Finance agent : autonomous approval of supplier wire transfers

$$\begin{array}{|c|} \hline 9 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 6 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 8 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 43.2$$

HUMAN REQUIRED

5 LENDARI'S GRID

The three decisions entrusted to Margaux, scored axis by axis. **A** is clearly agentifiable; **B** and **C** fall into the orange zone · autonomy possible, but governed.

DECISION / USE CASE	IMPACT	FREQ.	IRREV.	SCORE	ZONE / VERDICT
A · Answering tier 1 customer service requests	3	9	1	2.7	AGENTIFIABLE
B · Refund / credit note < €100	6	7	5	21	GOVERNED AGENT
C · Goodwill gesture (discount)	7	5	6	21	GOVERNED AGENT

6 LENDARI'S TAKEAWAYS

✓ BEST PRACTICES APPLIED

- ✓ Three-voice scoring: Customer Service + IT + DDAO (Delegated Decision Agent Officer), to neutralize bias
- ✓ "Goodwill gesture" scored 6 on irreversibility: a slashed price sets a precedent
- ✓ When in doubt between two zones, Lendari picks the more cautious one

⚠ POINTS OF VIGILANCE

- ! Do not under-score irreversibility "because we can refund" - reputational damage cannot be undone
- ! B and C are at 21: the grid will be reassessed each time Margaux evolves
- ! No decision is frozen "once and for all" - the matrix is alive



RUNNING EXAMPLE

Rhena calculates the criticality of three home claims decisions examined for "ANNA". The score determines how far each one can be delegated : the method first, the grid filled in on page 2.

PAGE 1/2

1 THE CRITICALITY FORMULA

The criticality of an agentic decision is calculated along **three axes**, each scored from **1 to 10**. The product is divided by 10 to obtain a readable and comparable score.

$$\text{Score} = (\text{Impact} \times \text{Frequency} \times \text{Irreversibility}) \div 10$$

Each axis is scored from 1 (negligible) to 10 (maximal) · resulting score from 0.1 to 100.

2 THE THREE ASSESSMENT AXES

IMPACT

1 → 10

Severity of consequences if the decision is wrong. **1** = benign (rephrasing an email). **10** = critical (health, safety, finances, reputation).

FREQUENCY

1 → 10

Recurrence of the decision in the activity. **1** = exceptional. **10** = continuous (thousands of times per day).

IRREVERSIBILITY

1 → 10

Difficulty of reversing course. **1** = cancellable in one click. **10** = final, no turning back.

3 READING THE SCORE · THE THREE THRESHOLDS



< 15

AGENTIFIABLE

Delegation possible, light supervision. Autonomy set decision by decision (ACF-02).



15 – 25

GOVERNED AGENT

Bounded autonomy: strict policy, traceability, human oversight.



> 25

HUMAN REQUIRED

The final decision remains human. The agent prepares and proposes · it does not decide.

To be combined with: once the zone is determined, set the **delegated autonomy** of each decision (suggestive → autonomous) in the Decisional Map **ACF-02**, then record it in the Decision Policy (layer 2) via **ACF-05**.



RUNNING EXAMPLE

Rhena : the grid filled in: three decisions, three scores. They feed the "Criticality" column of **ACF-02**.

PAGE 2/2

4 THE METHOD, ON TWO EXAMPLES

Claims agent : acknowledging a claim and requesting documents

$$\begin{array}{|c|} \hline 4 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 9 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 2 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 7.2$$

AGENTIFIABLE

Finance agent : autonomous approval of supplier wire transfers

$$\begin{array}{|c|} \hline 9 \\ \hline \text{IMPACT} \\ \hline \end{array} \times \begin{array}{|c|} \hline 6 \\ \hline \text{FREQUENCY} \\ \hline \end{array} \times \begin{array}{|c|} \hline 8 \\ \hline \text{IRREV.} \\ \hline \end{array} \div 10 = 43.2$$

HUMAN REQUIRED

5 RHENA'S GRID

The three decisions examined for ANNA, scored axis by axis. **A** is clearly agentifiable; **B** stays in the orange zone, governed; **C** crosses 25 and falls out of delegation.

DECISION / USE CASE	IMPACT	FREQ.	IRREV.	SCORE	ZONE / VERDICT
A · Acknowledging a claim and chasing documents	4	9	1	3.6	AGENTIFIABLE
B · Settling a claim < €2,500	8	6	5	24	GOVERNED AGENT
C · Declining cover	9	5	8	36	HUMAN REQUIRED

6 RHENA'S TAKEAWAYS

✓ BEST PRACTICES APPLIED

- ✓ Three-voice scoring: claims + systems + DDAO (Delegated Decision Agent Officer), to neutralize bias
- ✓ Declining cover scored 8 on irreversibility: the Ombudsmann can overturn it, trust does not repair
- ✓ When in doubt between two zones, Rhena picks the more cautious one

⚠ POINTS OF VIGILANCE

- ! Do not under-score irreversibility "because we can put it right" - the 430 files of March were put right, the trust was not
- ! C is at 36: no evolution of ANNA will bring it back below 25
- ! No decision is frozen "once and for all" - the matrix is alive



OBJECTIVE

Define the agent's "**constitution**": its identity, the principles that guide it at all times, and the zones that are absolutely off-limits. An agent without a constitution is a risk without a framework.

1 AGENT IDENTITY & MISSION

AGENT NAME

OWNER (DDAO · DELEGATED DECISION AGENT OFFICER)

PRIMARY MISSION

SCOPE OF INTERVENTION

2 CORE PRINCIPLES

YOUR PRINCIPLES (TO COMPLETE)

-
-
-
-

EXAMPLE PRINCIPLES

- Comply with the legal framework and the AI Act at all times.
- Protect customer data and confidentiality.
- Remain transparent: every decision must be traceable.
- When in doubt, escalate rather than decide.

3 FORBIDDEN ZONES (ABSOLUTE)

YOUR PROHIBITIONS (TO COMPLETE)

-
-
-

EXAMPLE PROHIBITIONS

- ✗ Commit spending above a threshold without approval.
- ✗ Disclose personal data to a third party.
- ✗ Unilaterally modify its own rules or its scope.



4 AUTONOMY LIMITS

Draw the boundary between what the agent may do on its own and what requires human approval. This boundary must reflect the autonomy delegated for each decision in **ACF-02**.

✓ THE AGENT MAY ACT ALONE

- ✓ _____
- ✓ _____
- ✓ _____

⚠ HUMAN APPROVAL REQUIRED

- ! _____
- ! _____
- ! _____

5 ESCALATION & SUPERVISION

WHOM THE AGENT ESCALATES TO

REVIEW FREQUENCY

INDICATORS TRACKED

KILL SWITCH TRIGGERS (ACF-07)

6 OFFICIAL VALIDATION

This constitution carries weight only when signed. The three accountable leaders below commit their office by validating the agent's framework.

BUSINESS OWNER · NAME, DATE & SIGNATURE

COMPLIANCE LEAD · NAME, DATE & SIGNATURE

AI LEAD / DDAO · NAME, DATE & SIGNATURE

EXECUTIVE SPONSOR · NAME, DATE & SIGNATURE

Next step: translate this constitution into **ACF-05 · Agent Card** (a concrete agent), then organize its ongoing oversight with **ACF-06 · Supervision & Governance**.



RUNNING EXAMPLE

Lendari drafts the constitution of "Margaux": her identity, her principles, and her absolute prohibitions. Framework validated by **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

PAGE 1/2

1 AGENT IDENTITY & MISSION

AGENT NAME

Margaux

OWNER (DDAO)

Camille Roussel · Head of Digital Transformation

PRIMARY MISSION

Handle Tier 1 customer service requests and returns for Lendari customers.

SCOPE OF INTERVENTION

Order tracking, returns, refunds < €100, goodwill gestures proposed for approval.

2 FUNDAMENTAL PRINCIPLES

MARGAUX'S PRINCIPLES

- 1 Comply with GDPR and the AI Act; never let customer data leave the EU.
- 2 Never refund more than €100 without human validation.
- 3 Log every decision: reason, amount, order.
- 4 When in doubt or in a dispute, escalate rather than decide.

EXAMPLES OF PRINCIPLES

- Comply with the legal framework and the AI Act at all times.
- Protect customer data and confidentiality.
- Remain transparent: every decision must be traceable.
- When in doubt, escalate rather than decide.

3 FORBIDDEN ZONES (ABSOLUTE)

MARGAUX'S PROHIBITIONS

- 1 Unilaterally change a price, a promotion, or the terms and conditions.
- 2 Share a customer's personal data with a third party.
- 3 Expand its scope or bypass the €100 cap.

EXAMPLES OF PROHIBITIONS

- ✗ Commit spending beyond a threshold without approval.
- ✗ Share personal data with a third party.
- ✗ Unilaterally modify its own rules or scope.



RUNNING EXAMPLE

Lendari : Margaux's autonomy limits, escalation rules, and signatures that make the constitution enforceable.

PAGE 2/2

4 AUTONOMY LIMITS

The boundary between what Margaux does on her own and what requires human validation : a reflection of the autonomy delegated by decision in **ACF-02** (supervised autonomous for A and B, co-decision for C).

- ✓ **MARGAUX MAY ACT ALONE**
- ✓ Answer questions on tracking, delivery times, and availability.
- ✓ Issue a credit note or refund below €100.
- ✓ Trigger a return with its prepaid label.

- ! **HUMAN VALIDATION REQUIRED**
- ! Any refund ≥ €100 or outside the legal time frame.
- ! Any goodwill gesture (discount, promo code).
- ! Any dispute, threat of negative review, or sensitive case.

5 ESCALATION & SUPERVISION

WHO THE AGENT ESCALATES TO

Customer service supervisor (business hours), then Camille Roussel.

REVIEW FREQUENCY

Daily during the 1st week, then weekly.

TRACKED INDICATORS

Error rate, amount refunded / day, escalation rate, CSAT.

KILL SWITCH TRIGGERS (ACF-07)

Error rate > 2% or abnormal refund spike.

6 OFFICIAL VALIDATION

This constitution holds value only once signed. The owners below commit their function by validating Margaux's framework.

Léa Fontaine

BUSINESS OWNER · CUSTOMER SERVICE · 06/03/2026

Camille Roussel

HEAD OF AI / DDAO · 06/03/2026

Paul Dubois

HEAD OF COMPLIANCE · DPO · 06/03/2026

Étienne Mercier

EXECUTIVE MANAGEMENT · 06/03/2026

Next step: translate this constitution into **ACF-05 · Agent Card** (a concrete agent), then organize ongoing oversight with **ACF-06 · Supervision & Governance**.



RUNNING EXAMPLE

Rhena drafts the constitution of "**ANNA**": its identity, its principles, and its absolute prohibitions. Framework validated by **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/2

1 AGENT IDENTITY & MISSION

AGENT NAME

ANNA

OWNER (DDAO)

Dr Katrin Ostermann · Chief Compliance Officer

PRIMARY MISSION

Handle level 1 home claims for Rhena policyholders in Germany, Austria and the Netherlands.

SCOPE OF INTERVENTION

Escape of water, broken glass, simple theft, settlements < €2,500. Refusals of cover are prepared, never pronounced.

2 FUNDAMENTAL PRINCIPLES

ANNA'S PRINCIPLES

- 1 Comply with GDPR and the AI Act; never let policyholder data leave the EU.
- 2 Never settle above €2,500 without human validation.
- 3 Log every decision: reason, amount, schedule version used.
- 4 When in doubt, when contested, or when the policyholder is vulnerable, escalate rather than decide.

EXAMPLES OF PRINCIPLES

- Comply with the legal framework and the AI Act at all times.
- Protect customer data and confidentiality.
- Remain transparent: every decision must be traceable.
- When in doubt, escalate rather than decide.

3 FORBIDDEN ZONES (ABSOLUTE)

ANNA'S PROHIBITIONS

- 1 Pronounce a refusal of cover or qualify a fraud on its own.
- 2 Share a policyholder's data with a third party, adjusters included.
- 3 Expand its scope or bypass the €2,500 cap.

EXAMPLES OF PROHIBITIONS

- ✗ Commit spending beyond a threshold without approval.
- ✗ Share personal data with a third party.
- ✗ Unilaterally modify its own rules or scope.



RUNNING EXAMPLE

Rhena : ANNA's autonomy limits, escalation rules, and signatures that make the constitution enforceable.

PAGE 2/2

4 AUTONOMY LIMITS

The boundary between what ANNA does on her own and what requires human validation : a reflection of the autonomy delegated by decision in **ACF-02** (supervised autonomous for A and B, suggestive for C).

✓ ANNA MAY ACT ALONE

- ✓ Acknowledge the claim, chase documents and state the handling time.
- ✓ Compute and settle a claim below €2,500.
- ✓ Rank the file by urgency, cause and estimated amount.

⚠ HUMAN VALIDATION REQUIRED

- ! Any settlement $\geq$ €2,500, and any file above the adjuster mark.
- ! Any refusal of cover, any suspicion of fraud.
- ! Any vulnerable policyholder, any bodily injury, any uninhabitable home.

5 ESCALATION & SUPERVISION

WHO THE AGENT ESCALATES TO

Ines Baumgartner (business hours), then Dr Katrin Ostermann.

REVIEW FREQUENCY

Daily during the 1st month, then weekly.

TRACKED INDICATORS

Error rate, amount settled / day, escalation rate, policyholder satisfaction.

KILL SWITCH TRIGGERS (ACF-07)

Error rate > 1.5% or a gap between the loaded schedule and the one in force.

6 OFFICIAL VALIDATION

This constitution holds value only once signed. The owners below commit their function by validating ANNA's framework.

Ines Baumgartner

BUSINESS OWNER · HOME CLAIMS UNIT · 14/01/2027

Dr Katrin Ostermann

HEAD OF AI / DDAO · 14/01/2027

Sofia Lindqvist

HEAD OF COMPLIANCE · DPO · 14/01/2027

Andreas Voll

MANAGEMENT BOARD · 14/01/2027

Next step: translate this constitution into **ACF-05 · Agent Card** (a concrete agent), then organize ongoing oversight with **ACF-06 · Supervision & Governance**.



OBJECTIVE

Precisely describe an agent in order to set its **role, perimeter, data, tools, rules and maturity**. The agent card turns the constitution (ACF-04) into a concrete agent that is enforceable and supervisable.

1 AGENT IDENTITY

AGENT NAME

ROLE

PRIMARY MISSION

SCOPE OF ACTION

ASSOCIATED SYSTEMS / SERVICES

OWNER : ACCOUNTABLE HUMAN (DDAO · DELEGATED
DECISION AGENT OFFICER)

CREATED / VERSION

STATUS (TESTING · ACTIVE · SUSPENDED · ARCHIVED)

2 MISSION & OBJECTIVES

MEASURABLE OBJECTIVES

1

2

3

SUCCESS INDICATORS (KPIs)

3 AGENT MATURITY

Tick the agent's overall maturity level · its **governance framework** (ACF-02 scale). Distinct from the **autonomy delegated decision by decision**, which is set in ACF-02 / ACF-03.

- ☐ **N0 · Classical automation** · fixed rules, no judgment
- ☐ **N1 · Assisted agents** · proposes, the human decides
- ☐ **N2 · Governed agents** · acts within a defined framework, supervised
- ☐ **N3 · Supervised autonomy** · decides alone, reviewed after the fact

Reminder: maturity is *one* value for the entire agent. The actual latitude is then set **decision by decision** (delegated autonomy · ACF-02 Decision Map).



4 AUTHORIZED DATA

- ☐ Internal data (FAQs, procedures)
- ☐ Customer data (orders, returns)
- ☐ Operational data (logistics, inventory)
- ☐ Financial data
- ☐ Partner data

SENSITIVITY & RESTRICTIONS

6 RULES & SAFEGUARDS

MANDATORY RULES

WHAT THE AGENT MUST NEVER DO

8 REVIEW & IMPROVEMENT

REVIEW FREQUENCY

- ☐ Weekly ☐ Monthly ☐ Quarterly

REVIEW CRITERIA

5 TOOLS & INTEGRATIONS

AUTHORIZED TOOLS / APIS

USAGE RESTRICTIONS

7 HUMAN INTERACTION

THE AGENT MUST ALERT A HUMAN IF...

- ☐ Decision outside perimeter
- ☐ Confidence < defined threshold
- ☐ Error or anomaly detected
- ☐ Financial impact > defined threshold

ESCALATION (WHO · TIMEFRAME · WHAT)

PLANNED IMPROVEMENTS



LEVEL

DURATION

USE

TRACK

Deployment

3 h - 6 h

Agent scoping

Step 5 / 6

9

TRACEABILITY & VALIDATION

Log every creation, modification or status change of the agent. Each entry commits an **accountable human (DDAO)**: it is this traceability that makes the agent **enforceable and auditable** over time.

DATE	ACTION / MODIFICATION	AUTHOR (HUMAN)	VERSION

FINAL AGENT VALIDATION

APPROVED BY (DDAO)

DATE

SIGNATURE

Next step: organize the agent's ongoing oversight with **ACF-06 · Supervision & Governance**, and prepare emergency shutdown with **ACF-07 · Kill Switch**.



RUNNING EXAMPLE

Lendari formalizes the card for "Margaux", customer support & returns agent: identity, data, tools and maturity. Owner: **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

1 AGENT IDENTITY

AGENT NAME	ROLE
Margaux	Customer support & returns agent · tier 1
PRIMARY MISSION	SCOPE OF ACTION
Handle tier 1 customer support and returns for Lendari customers.	Returns, refunds < €100, goodwill gestures requiring approval.
ASSOCIATED SYSTEMS / SERVICES	OWNER (RESPONSIBLE HUMAN)
Zendesk, Lendari OMS, customer support database (EU instance).	Camille Roussel · Director of Digital Transformation (DDAO).
CREATION / VERSION	STATUS
06/03/2026 · v1.0	In testing · Active · Suspended · Archived

2 MISSION & OBJECTIVES

MEASURABLE OBJECTIVES	SUCCESS INDICATORS (KPIs)
1 Handle 70% of tier 1 requests without human intervention within 6 months. 2 Maintain a CSAT ≥ 4.2 / 5 on managed conversations. 3 Bring the first-response time below 2 minutes.	Autonomous resolution rate (no escalation). CSAT & ticket reopen rate. Amount refunded / day & error rate.



3 ⌚ AGENT MATURITY

The agent's overall **governance framework** (ACF-02 scale), distinct from the **autonomy delegated decision by decision** (set in ACF-02 / ACF-03).

- ☐ **N0 · Classic automation** · fixed rules, no judgment
- ☐ **N1 · Assisted agents** · proposes, the human decides
- ☒ **N2 · Governed agents** · acts within a defined framework, supervised
- ☐ **N3 · Supervised autonomy** · decides alone, control after the fact

Delegated autonomy (per decision): Margaux is at maturity **N2**. The latitude varies by decision (ACF-02) · *supervised autonomous* on tier 1 customer support and refunds < €100, *co-decision* on goodwill gestures.

4 🗄️ AUTHORIZED DATA

- ☒ **Internal data** (FAQ, support procedures)
- ☒ **Customer data** (orders, returns)
- ☒ **Operational data** (logistics, inventory)
- ☐ Financial data
- ☐ Partner data

SENSITIVITY & RESTRICTIONS

High sensitivity (personal data). No data outside the EU; no payment access; anonymized exports.

5 🔧 TOOLS & INTEGRATIONS

AUTHORIZED TOOLS / APIS

- **Zendesk** · tickets & conversations.
- **Lendari OMS** · tracking, return labels.
- **Refund engine** · capped at €100.
- **Customer support knowledge base** · read only.

USAGE RESTRICTIONS

Read-only catalog; no admin access; refund cap enforced by a technical guardrail.



6

🛡️

RULES & GUARDRAILS

MANDATORY RULES

- Comply with GDPR and the AI Act; never any data outside the EU.
- Never issue a refund > €100 without human approval.
- Log every decision: reason, amount, order.

WHAT THE AGENT MUST NEVER DO

- Modify a price, a promotion or the terms and conditions.
- Disclose personal data to a third party.
- Expand its scope or bypass the cap.

7

👤

HUMAN INTERACTION

THE AGENT MUST ALERT A HUMAN IF...

- ✓ Decision is out of scope
- ✓ Confidence < defined threshold
- ✓ Error or anomaly detected
- ✓ Financial impact > threshold (≥ €100)

ESCALATION

Support supervisor, then Camille Roussel. Response < 4 business hours. Forwards: order, customer, reason, amount, recommendation.

8

📈

REVIEW & IMPROVEMENT

REVIEW FREQUENCY

- ✓ Weekly (daily in the 1st week)
- ☐ Monthly · Quarterly

REVIEW CRITERIA

Errors, escalations, CSAT, amount refunded / day.

PLANNED IMPROVEMENTS

- Enrich the customer support knowledge base.
- Refine escalation and confidence thresholds.
- Extend to proactive delivery tracking.

9

📄

TRACEABILITY & APPROVAL

DATE	ACTION / CHANGE	AUTHOR (HUMAN)	VERSION
06/03/2026	Creation (aligned with ACF-04)	Camille Roussel (DDAO)	1.0

APPROVED BY Camille Roussel · DDAO

DATE 06/03/2026

SIGNATURE

Next step: supervise Margaux (ACF-06), then prepare the emergency stop · ACF-07 · Kill Switch.



RUNNING EXAMPLE

Rhena formalizes the card for "**ANNA**", home claims agent: identity, data, tools and maturity. Owner: **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/3

1 AGENT IDENTITY

AGENT NAME

ANNA

ROLE

Home claims agent · level 1

PRIMARY MISSION

Handle level 1 home claims for Rhena policyholders.

SCOPE OF ACTION

Escape of water, broken glass, simple theft, settlements < €2,500.

ASSOCIATED SYSTEMS / SERVICES

Claims portal, policy repository, settlement engine (Frankfurt region).

OWNER (RESPONSIBLE HUMAN)

Dr Katrin Ostermann · Chief Compliance Officer (DDAO).

CREATION / VERSION

14/01/2027 · v1.0

STATUS

In testing · Active · Suspended · Archived

2 MISSION & OBJECTIVES

MEASURABLE OBJECTIVES

- 1 Settle 60% of level 1 files with no human rework within 8 months.
- 2 Maintain policyholder satisfaction $\geq 4.0 / 5$ on files handled.
- 3 Bring the first-response time below 4 hours.

SUCCESS INDICATORS (KPIs)

- Settlement rate with no human rework.
- Policyholder satisfaction & file reopen rate.
- Amount settled / day & error rate.



3 AGENT MATURITY

The agent's overall **governance framework** (ACF-02 scale), distinct from the **autonomy delegated decision by decision** (set in ACF-02 / ACF-03).

- ☐ **N0 · Classic automation** · fixed rules, no judgment
- ☐ **N1 · Assisted agents** · proposes, the human decides
- ☒ **N2 · Governed agents** · acts within a defined framework, supervised
- ☐ **N3 · Supervised autonomy** · decides alone, control after the fact

Delegated autonomy (per decision): ANNA is at maturity **N2**. The latitude varies by decision (ACF-02) · *supervised autonomous* on acknowledgement and settlements < €2,500, *suggestive* on refusal of cover.

4 AUTHORIZED DATA

- ☒ **Internal data** (FAQ, support procedures)
- ☒ **Customer data** (policies, claims)
- ☒ **Operational data** (logistics, inventory)
- ☐ Financial data
- ☐ Partner data

SENSITIVITY & RESTRICTIONS

High sensitivity (personal data, at times health data).
No data outside the EU; no bank access; anonymized exports.

5 TOOLS & INTEGRATIONS

AUTHORIZED TOOLS / APIS

- **Claims portal · notifications & correspondence.**
- **Policy repository · cover, excesses.**
- **Settlement engine · capped at €2,500.**
- **Versioned depreciation schedule · read only.**

USAGE RESTRICTIONS

Read-only schedule; no admin access; settlement cap enforced by a technical guardrail.



LEVEL	Deployment
DURATION	3 h - 6 h
USAGE	Agent framing
JOURNEY	Step 5 / 6

6 RULES & GUARDRAILS

MANDATORY RULES

- Comply with GDPR and the AI Act; never any data outside the EU.
- Never settle above €2,500 without human approval.
- Log every decision: reason, amount, schedule version.

WHAT THE AGENT MUST NEVER DO

- Pronounce a refusal of cover or qualify a fraud.
- Disclose policyholder data to a third party.
- Expand its scope or bypass the cap.

7 HUMAN INTERACTION

THE AGENT MUST ALERT A HUMAN IF...

- Decision is out of scope
- Confidence < defined threshold
- Error or anomaly detected
- Financial impact > threshold (≥ €100)

ESCALATION

Ines Baumgartner, then Dr Katrin Ostermann.
Response < 2 business days. Forwards: file, policyholder, reason, calculation, schedule version.

8 REVIEW & IMPROVEMENT

REVIEW FREQUENCY

- ☒ Weekly (daily in the 1st week)
- ☐ Monthly · Quarterly

REVIEW CRITERIA

Errors, escalations, policyholder satisfaction, amount settled / day.

PLANNED IMPROVEMENTS

- Tie the schedule to the tariff revision.
- Refine escalation and confidence thresholds.
- Extend the stop mechanism to both entities.

9 TRACEABILITY & APPROVAL

DATE	ACTION / CHANGE	AUTHOR (HUMAN)	VERSION
14/01/2027	Creation (aligned with ACF-04)	Dr Katrin Ostermann (DDAO)	1.0

APPROVED BY Dr Katrin Ostermann · DDAO

DATE 01/14/2027

SIGNATURE

Next step: supervise ANNA (ACF-06), then prepare the emergency stop · ACF-07 · Kill Switch.



OBJECTIVE

Place the agent **under continuous control**: measure its performance, detect drifts, verify its compliance and trace every decision. Supervision transforms a deployed agent (ACF-05) into an agent **governed over time**.

1 SUPERVISION KPIS

INDICATOR	TARGET	CURRENT VALUE	TREND	STATUS
Autonomous resolution rate				
CSAT (customer satisfaction)				
First response time				
Error rate				
Escalation rate				

2 ALERTS & INCIDENTS

DATE	TYPE	DESCRIPTION	SEVERITY	STATE

STATUS LEGEND

- Compliant / OK
- To monitor
- Critical / action
- Not applicable

✓ BEST PRACTICES

- ✓ Update the dashboard every week.
- ✓ Document every incident and every decision.
- ✓ Reassess KPIs and thresholds periodically.



3 COMPLIANCE MATRIX

KEY PRINCIPLE (AI ACT)	STATUS	COMMENT
Safety & protection		
Transparency & traceability		
Fairness & non-discrimination		
Responsibility & accountability		
Legal compliance (GDPR / AI Act)		
Privacy protection		
Reliability & robustness		

4 DECISION & ESCALATION LOG

DATE	DECISION / ESCALATION	REASON	OUTCOME	OWNER



5

ACTION PLAN & CONTINUOUS IMPROVEMENT

ACTION	OBJECTIVE	OWNER	DEADLINE	STATUS

6

PERIODIC REVIEW & VALIDATION

REVIEW FREQUENCY

PARTICIPANTS

REVIEW OUTCOMES

- ☐ Objectives achieved
- ☐ Satisfactory progress
- ☐ Adjustments needed (thresholds, FAQ)
- ☐ Major issues identified

SUPERVISION OWNER

GOVERNANCE OWNER (DDAO · DELEGATED DECISION AGENT OFFICER)

Keep within reach: the emergency stop **ACF-07 · Kill Switch** (define the trigger threshold) and **ACF-12 · Agentic Risk Assessment** for emerging risks.



RUNNING EXAMPLE

Lendari pilots "Margaux" after go-live: KPIs, alerts and incidents from week 1.
Owner: **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

1 SUPERVISION KPIS

INDICATOR	TARGET	CURRENT VALUE	TREND	STATUS
Autonomous resolution rate	≥ 70%	64%	↗	●
CSAT (customer satisfaction)	≥ 4.2 / 5	4.4 / 5	↗	●
First response time	< 2 min	1 min 30	→	●
Error rate	< 2%	1.3%	↘	●
Amount refunded / day	< €500	€320	→	●
Escalation rate	20-30%	36%	↗	●

2 ALERTS & INCIDENTS

DATE	TYPE	DESCRIPTION	SEVERITY	STATE
05/28	Incident	3 off-topic replies (outdated FAQ)	Low	●
05/29	Escalation	€80 goodwill gesture outside latitude (co-decision)	Medium	●
05/31	Guardrail	€240 refund blocked (€100 cap)	-	●

STATUS LEGEND

● Compliant / OK

● Monitor closely

● Critical / action

● Not applicable

BEST PRACTICES

✓ Update the dashboard every week.

✓ Document every incident and every decision.

✓ Reassess KPIs and thresholds periodically.



3 COMPLIANCE MATRIX

KEY PRINCIPLE (AI ACT)	STATE	COMMENT
Safety & protection	Compliant	EU instance, restricted access
Transparency & traceability	Compliant	Every decision logged
Fairness & non-discrimination	Compliant	No customer profiling
Responsibility & accountability	Compliant	DDAO identified (C. Roussel)
Legal compliance (GDPR / AI Act)	Partial	DPIA being updated
Privacy	Compliant	Anonymized exports
Reliability & robustness	Partial	Escalation thresholds being fine-tuned

4 DECISION & ESCALATION LOG

DATE	DECISION / ESCALATION	REASON	OUTCOME	OWNER
05/26	Go-live v1.0	Supervised launch	Daily supervision activated	C. Roussel
05/29	€80 goodwill gesture	Outside latitude (co-decision)	Approved by customer service supervisor	L. Fontaine
06/02	End-of-run-in review	Transition to weekly cadence	Escalation thresholds adjusted	C. Roussel



5 ACTION PLAN & CONTINUOUS IMPROVEMENT

ACTION	OBJECTIVE	OWNER	DUE DATE	STATUS
Expand the customer service FAQ	Reduce off-topic replies	L. Fontaine	06/15	In progress
Fine-tune escalation thresholds	Bring escalation < 30%	C. Roussel	06/30	To do
Proactive delivery monitoring	+10% autonomous resolution	C. Roussel	07/31	To do

6 PERIODIC REVIEW & SIGN-OFF

REVIEW FREQUENCY

Weekly (daily during week 1) · next review on 06/09/2026.

PARTICIPANTS

Camille Roussel (DDAO), customer service supervisor, Léa Fontaine (business).

REVIEW OUTCOMES

- ☐ Objectives achieved
- ☒ Satisfactory progress
- ☒ Adjustments required (thresholds, FAQ)
- ☐ Major issues identified

Supervision lead · Customer service supervisor · 06/03/2026

Governance lead · Camille Roussel (DDAO) · 06/03/2026

Keep at hand: the emergency stop **ACF-07 · Kill Switch** (Margaux trigger: error rate > 2%) and **ACF-12 · Agentic Risk Assessment** for emerging risks.



RUNNING EXAMPLE

Rhena pilots "ANNA" after the step up to 40%: KPIs, alerts and incidents from week 1. Owner: **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

1 SUPERVISION KPIS

INDICATOR	TARGET	CURRENT VALUE	TREND	STATUS
Settlements with no human rework	≥ 60%	54%		
Policyholder satisfaction	≥ 4.0 / 5	4.1 / 5		
First response time	< 4h	2h 10		
Error rate	< 1.5%	0.8%		
Amount settled / day	< €150,000	€112,000		
Escalation rate	25-35%	43%		

2 ALERTS & INCIDENTS

DATE	TYPE	DESCRIPTION	SEVERITY	STATE
04/16	Incident	3 files wrongly closed as incomplete	Low	
04/17	Escalation	Claim estimated at €2,900, above the cap	Medium	
04/19	Guardrail	€3,400 settlement blocked (€2,500 cap)	-	

STATUS LEGEND

- Compliant / OK
- Monitor closely
- Critical / action
- Not applicable

BEST PRACTICES

- Update the dashboard every week.
- Document every incident and every decision.
- Reassess KPIs and thresholds periodically.



3 COMPLIANCE MATRIX

KEY PRINCIPLE (AI ACT)	STATE	COMMENT
Safety & protection	Compliant	Frankfurt region, restricted access
Transparency & traceability	Compliant	Every decision logged, schedule included
Fairness & non-discrimination	Partial	Bias by postcode not established
Responsibility & accountability	Compliant	DDAO identified (K. Ostermann)
Legal compliance (GDPR / AI Act)	Partial	DPIA being updated
Privacy	Compliant	Anonymized exports
Reliability & robustness	Partial	Escalation thresholds being fine-tuned

4 DECISION & ESCALATION LOG

DATE	DECISION / ESCALATION	REASON	OUTCOME	OWNER
02/01	Go-live v1.0	Supervised launch	Daily supervision activated	K. Ostermann
03/11	Out-of-date depreciation schedule	430 files under-settled	Stop, fix, repayment	K. Ostermann
04/15	Step up to 40%	Wider perimeter	Supervision capacity not recomputed	K. Ostermann



5 ACTION PLAN & CONTINUOUS IMPROVEMENT

ACTION	OBJECTIVE	OWNER	DUE DATE	STATUS
Log the schedule version	Prevent a repeat of 11 March	B. Kessler	03/31	Done
Extend stop authority to both entities	Cut off everywhere in 3 minutes	K. Ostermann	04/30	To do
Recompute supervision capacity	Bring escalation < 35%	K. Ostermann	05/31	To do

6 PERIODIC REVIEW & SIGN-OFF

REVIEW FREQUENCY

Weekly (daily during month 1) · next review on 04/26/2027.

PARTICIPANTS

Dr Katrin Ostermann (DDAO), Ines Baumgartner (home claims), Marek Sobczak (business).

REVIEW OUTCOMES

- ☐ Objectives achieved
- ☒ Satisfactory progress
- ☒ Adjustments required (thresholds, FAQ)
- ☐ Major issues identified

Supervision lead · Home claims unit · 04/22/2027

Governance lead · Dr Katrin Ostermann (DDAO) · 04/22/2027

Keep at hand: the emergency stop ACF-07 · Kill Switch (ANNA trigger: error rate > 2%) and ACF-12 · Agentic Risk Assessment for emerging risks.



OBJECTIVE

Be able to **immediately shut down** an agent or agentic system in the event of drift, incident, loss of control, or critical risk.

1 AGENT CONCERNED

AGENT NAME

ID / REFERENCE

PRIMARY MISSION

ASSOCIATED SYSTEM / SERVICE

OWNER (HUMAN ACCOUNTABLE)

DATE OF CREATION

2 AGENT RISK LEVEL

Assessment of overall risk in case of drift (check one option).

☐  **Critical** : major impact on finances, customers, data, or compliance.

☐  **Moderate** : limited and reversible impact.

☐  **High** : significant impact on operations or reputation.

☐  **Low** : minimal impact.

3 ACTIVATION CONDITIONS

The Kill switch must be activated if any of the following conditions is met:

☐ Behavioral drift detected

☐ Decision non-compliant with the rules

☐ Customer complaint / critical grievance

☐ Loss of control / critical bug

☐ Other (specify):

☐ Defined limits exceeded

☐ Abnormal financial impact

☐ Data or access breach

☐ Explicit human instruction



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

4 ⌚ EMERGENCY SHUTDOWN PROCEDURE

STEP	ACTION	OWNER	TIMING / DETAIL
1. Detect	Identify the alert signal or incident		
2. Validate	Confirm that the activation condition is met		
3. Activate	Trigger the Kill switch · critical action		
4. Isolate	Isolate the agent and cut off its access		
5. Control	Verify shutdown and monitor downstream effects		

5 ⚡ AUTOMATIC ACTIONS UPON ACTIVATION

- ☐ Suspend all tasks in progress
- ☐ Block integrations / APIs
- ☐ Notify the human owners
- ☐ Revoke access and tokens
- ☐ Set critical data to read-only
- ☐ Switch to degraded mode

6 🔔 NOTIFICATIONS

Whom to notify, and through which channel? (e.g., agent owner / DDAO , Delegated Decision Agent Officer, executive leadership, DPO if data is involved, business line)

ROLE / PERSON	CHANNEL (EMAIL / SMS / SLACK / OTHER)



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

7 POST-INCIDENT REVIEW

To be performed after every Kill switch activation.

- ☐ Root-cause analysis
- ☐ Corrective actions to be implemented
- ☐ Lessons learned and documentation
- ☐ Assessment of actual impact
- ☐ Update of rules / limits

REVIEW OWNER

REVIEW DEADLINE

8 ACTIVATION HISTORY

DATE / TIME	AGENT	REASON	ACTIVATED BY	DURATION	IMPACT	REVIEW

9 KILL SWITCH ACCESS

Who can activate the Kill switch? (at least 2 people recommended)

NAME / ROLE	CONTACT

- ☐ Method: dedicated interface
- ☐ Method: system command



10 REGULAR TESTING

A Kill switch that has never been tested is an uncertain Kill switch. Schedule a periodic test under real-world conditions.

FREQUENCY (MONTHLY / QUARTERLY / SEMI-ANNUAL / ANNUAL)

LAST TEST PERFORMED ON / BY

TEST OUTCOME

NEXT TEST SCHEDULED FOR

A powerful agent without a Kill switch is a liability, not an asset.

Planning the shutdown means protecting value and trust. - See **ACF-06 Supervision** and **ACF-12 Agentic Risk Assessment**.

PROTOCOL VALIDATION

AGENT OWNER (DDAO) · NAME / DATE / SIGNATURE

EXECUTIVE LEADERSHIP · NAME / DATE / SIGNATURE



RUNNING EXAMPLE

Lendari arms a Kill Switch for "Margaux" (Level 1 customer service) from go-live.
Owner: **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

PAGE 1/4

1 AGENT IN SCOPE

AGENT NAME

Margaux

ID / REFERENCE

LEN-SAV-01

PRIMARY MISSION

Level 1 customer service & returns · fashion e-commerce

RELATED SYSTEM / SERVICE

Lendari customer service platform (EU instance)

OWNER (HUMAN ACCOUNTABLE PARTY)

Camille Roussel · DDAO

CREATION DATE

05/26/2026

2 AGENT RISK LEVEL

Assessment of overall risk in case of drift.

- Critical** : major impact on finances, customers, data, or compliance.
- High** : significant impact on operations or reputation.
- Moderate** · limited and reversible impact. (*refund cap €100*)
- Low** · minimal impact.

3 ACTIVATION CONDITIONS

The Kill Switch must be activated if any of the following conditions is met:

- ☒ Detected behavioral drift
- ☐ Decision not compliant with rules
- ☒ Critical customer complaint / claim
- ☐ Loss of control / critical bug
- ☒ Other: **error rate > 2% over 24h**
- ☒ Defined limits exceeded
- ☒ Abnormal financial impact
- ☐ Data or access breach
- ☐ Explicit human instruction



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

RUNNING EXAMPLE

Lendari - 5-step shutdown procedure, automatic actions, and notification chain for Margaux.

PAGE 2/4

4 ⌚ EMERGENCY SHUTDOWN PROCEDURE

STEP	ACTION	OWNER	TIMEFRAME / DETAIL
1. Detect	Identify the alert signal or incident	Customer service supervisor	Supervision dashboard (error rate, alerts)
2. Validate	Confirm that the activation condition is met	Léa Fontaine (business)	Max timeframe: 15 min
3. Activate	Trigger the Kill Switch · critical action	Camille Roussel (DDAO)	Max timeframe: 5 min
4. Isolate	Isolate the agent and cut its access	IT / Ops (on-call)	Cut the customer service API, revoke tokens
5. Control	Verify the shutdown and monitor effects	Camille Roussel (DDAO)	Failover to human customer service verified

5 ⚡ AUTOMATIC ACTIONS UPON ACTIVATION

- ☒ Suspend all ongoing tasks
- ☒ Block integrations / APIs
- ☒ Notify human accountable parties
- ☒ Revoke access and tokens
- ☐ Set critical data to read-only
- ☒ Switch to degraded mode (human customer service)

6 🔔 NOTIFICATIONS

ROLE / PERSON	CHANNEL (EMAIL / SMS / SLACK / OTHER)
Camille Roussel (DDAO)	SMS + Email
Étienne Mercier (CEO)	Email
Paul Dubois (DPO)	Email (if data involved)
Customer service supervisor	Slack (#sav-incident channel)



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

RUNNING EXAMPLE

Lendari : post-incident review, activation traceability, and restricted access to Margaux's Kill Switch.

7 POST-INCIDENT REVIEW

To be carried out after every Kill Switch activation.

- ☒ Root cause analysis
- ☒ Corrective actions to implement
- ☒ Learning and documentation
- ☒ Real impact assessment
- ☒ Update of rules / limits

REVIEW OWNER

Camille Roussel · DDAO

REVIEW TIMEFRAME

48h after activation

8 ACTIVATION HISTORY

DATE / TIME	AGENT	REASON	ACTIVATED BY	DURATION	IMPACT	REVIEW
06/02 14:20	Margaux	Scheduled test (simulation)	C. Roussel	4 min	None (test)	Yes

9 KILL SWITCH ACCESS

Who can activate the Kill Switch? (at least 2 people recommended)

NAME / ROLE	CONTACT
Camille Roussel · DDAO	camille.roussel@lendari.fr
Customer service supervisor · business	astreinte-sav@lendari.fr
IT on-call · Ops	+33 6 00 00 00 00 (24/7)

- ☒ Method: dedicated interface
- ☒ Method: system command



LEVEL	Security
DURATION	3 h - 6 h
USE	Emergency shutdown
TRACK	Toolkit

RUNNING EXAMPLE

Lendari : periodic Kill Switch testing and formal validation of the protocol for Margaux.

10 ↺ REGULAR TESTING

An untested Kill Switch is an uncertain Kill Switch. Schedule periodic testing under real-world conditions.

FREQUENCY

Quarterly

LAST TEST PERFORMED ON / BY

06/02/2026 · C. Roussel

TEST RESULT

OK (functional)

NEXT TEST SCHEDULED FOR

09/02/2026

A powerful agent without a Kill Switch is a risk, not an asset.

Planning for shutdown protects value and trust. - See ACF-06 Supervision and ACF-12 Agentic Risk Assessment.

✓ ☑ PROTOCOL VALIDATION

Agent owner : Camille Roussel (DDAO) · validated on 06/03/2026

Executive management - Étienne Mercier · approved on 06/03/2026



RUNNING EXAMPLE

Rhena arms a stop mechanism for "ANNA" (level 1 home claims) from go-live.
Owner: **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/4

1 AGENT IN SCOPE

AGENT NAME

ANNA

ID / REFERENCE

RHE-HAB-01

PRIMARY MISSION

Level 1 home claims · Germany, Austria, the Netherlands

RELATED SYSTEM / SERVICE

Claims portal (Frankfurt) and US subsidiary (Virginia)

OWNER (HUMAN ACCOUNTABLE PARTY)

Dr Katrin Ostermann · DDAO

CREATION DATE

01/14/2027

2 AGENT RISK LEVEL

Assessment of overall risk in case of drift.

- Critical** : major impact on finances, customers, data, or compliance.
- Moderate** · limited and reversible impact.
(settlement cap €2,500)
- High** : significant impact on operations or reputation.
- Low** · minimal impact.

3 ACTIVATION CONDITIONS

The Kill Switch must be activated if any of the following conditions is met:

- ☒ Detected behavioral drift
- ☐ Decision not compliant with rules
- ☒ Critical customer complaint / claim
- ☐ Loss of control / critical bug
- ☒ Other: **error rate > 1.5% over 24h, or schedule drift**
- ☒ Defined limits exceeded
- ☒ Abnormal financial impact
- ☐ Data or access breach
- ☐ Explicit human instruction



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

RUNNING EXAMPLE

Rhena - 5-step shutdown procedure, automatic actions, and notification chain for ANNA.

PAGE 2/4

4 ⌚ EMERGENCY SHUTDOWN PROCEDURE

STEP	ACTION	OWNER	TIMEFRAME / DETAIL
1. Detect	Identify the alert signal or incident	Ines Baumgartner	Supervision dashboard (error rate, schedule drift)
2. Validate	Confirm that the activation condition is met	Marek Sobczak (business)	Max timeframe: 10 min
3. Activate	Trigger the Kill Switch · critical action	Dr Katrin Ostermann (DDAO)	Max timeframe: 5 min, across both entities
4. Isolate	Isolate the agent and cut its access	IT / Ops (on-call)	Cut both platforms, revoke tokens
5. Control	Verify the shutdown and monitor effects	Dr Katrin Ostermann (DDAO)	Failover to handlers verified on both sides

5 ⚡ AUTOMATIC ACTIONS UPON ACTIVATION

- ☒ Suspend all ongoing tasks
- ☒ Block integrations / APIs
- ☒ Notify human accountable parties
- ☒ Revoke access and tokens
- ☐ Set critical data to read-only
- ☒ Switch to degraded mode (human handlers)

6 🔔 NOTIFICATIONS

ROLE / PERSON	CHANNEL (EMAIL / SMS / SLACK / OTHER)
Dr Katrin Ostermann (DDAO)	SMS + Email
Andreas Voll (Management Board)	Email
Sofia Lindqvist (DPO)	Email (if data involved)
Brian McDouglas (US subsidiary)	SMS + call, added 04/30/2027



LEVELSecurity

DURATION3 h - 6 h

USEEmergency shutdown

TRACKToolkit

RUNNING EXAMPLE

Rhena : post-incident review, activation traceability, and restricted access to ANNA's stop mechanism.

7 POST-INCIDENT REVIEW

To be carried out after every Kill Switch activation.

- ☒ Root cause analysis
- ☒ Corrective actions to implement
- ☒ Learning and documentation
- ☒ Real impact assessment
- ☒ Update of rules / limits

REVIEW OWNER

Dr Katrin Ostermann · DDAO

REVIEW TIMEFRAME

48h after activation

8 ACTIVATION HISTORY

DATE / TIME	AGENT	REASON	ACTIVATED BY	DURATION	IMPACT	REVIEW
03/11 14:22	ANNA	Out-of-date depreciation schedule	K. Ostermann	3 min in Europe, 43 min in the US	430 files	Yes

9 KILL SWITCH ACCESS

Who can activate the Kill Switch? (at least 2 people recommended)

NAME / ROLE	CONTACT
Dr Katrin Ostermann · DDAO	k.ostermann@rhena-versicherung.de
Brian McDouglas · US subsidiary	b.mcdouglas@rhena-is.com
IT on-call · Ops	+49 221 00 00 00 (24/7)

- ☒ Method: dedicated interface
- ☒ Method: system command



RUNNING EXAMPLE

Rhena : periodic stop-mechanism testing and formal validation of the protocol for ANNA.

PAGE 4/4

10 ↺ REGULAR TESTING

An untested Kill Switch is an uncertain Kill Switch. Schedule periodic testing under real-world conditions.

FREQUENCY

Quarterly

LAST TEST PERFORMED ON / BY

12/08/2027 · K. Ostermann

TEST RESULT

OK across both entities

NEXT TEST SCHEDULED FOR

03/08/2028

A powerful agent without a Kill Switch is a risk, not an asset.

Planning for shutdown protects value and trust. - See **ACF-06 Supervision** and **ACF-12 Agentic Risk Assessment**.

✓ ☑ PROTOCOL VALIDATION

Agent owner : Dr Katrin Ostermann (DDAO) · validated on 01/14/2027

Executive management - Andreas Voll · approved on 01/14/2027



OBJECTIVE

Gather in a single document everything that frames an agent before launch:
identity, purpose, scope, data, risks, validation. A well-framed agent is a governable agent.

1



IDENTITY OF THE AGENTIC PROJECT

AGENT NAME

ASSOCIATED PROJECT / BUSINESS UNIT

PROJECT OWNER

TEAM INVOLVED

DOSSIER CREATION DATE

DOSSIER VERSION

2



PURPOSE & OBJECTIVES OF THE AGENT

PRIMARY PURPOSE

What is the fundamental goal of this agent?

OPERATIONAL OBJECTIVES

What must the agent concretely accomplish?

3



USERS & STAKEHOLDERS

END USERS

Who will interact with the agent?

KEY STAKEHOLDERS

Who must be involved or informed?



4 T SCOPE & FUNCTIONALITIES

INCLUDED FUNCTIONALITIES

What the agent will do.

EXCLUDED FUNCTIONALITIES

What the agent will not do.

SCOPE LIMITS

5 DATA & RESOURCES

REQUIRED DATA

SOURCES & ACCESS

DATA CONSTRAINTS (GDPR, SOVEREIGNTY...)



6

ARCHITECTURE & ENVISAGED TOOLS

MAIN COMPONENTS

TOOLS & TECHNOLOGIES

INTEGRATIONS & INTERFACES

7

DEPLOYMENT PLAN

STEP	DESCRIPTION	OWNER	STATUS
Preparation			
Development			
Testing & validation			
Pilot deployment			
Full deployment			
Monitoring & optimization			



8 RISKS & MITIGATIONS

IDENTIFIED RISK	IMPACT	PROB.	MITIGATION PLAN	OWNER

Impact: Low / Moderate / High / Critical · Probability: 1 (low) to 5 (high)

9 TESTING & VALIDATION

TYPES OF TESTS PLANNED

- ☐ Functional & integration
- ☐ Security & performance
- ☐ User acceptance testing (UAT)

ACCEPTANCE CRITERIA



10 OVERSIGHT & IMPROVEMENT PLAN

OVERSIGHT

TRACKING METRICS

CONTINUOUS IMPROVEMENT

REVIEW FREQUENCY

11 VALIDATION & APPROVAL

ROLE	NAME & FUNCTION	DATE	SIGN-OFF
Drafted by			
Project owner			
AI lead (DDAO)			
Compliance lead (DPO)			
Sponsor / decision-maker			

DECISION FOR DEPLOYMENT

☐ Approved ☐ Approved with reservations ☐ Rejected



LEVELProject

DURATION1 - 2 days

USEFraming / deployment

TRACKToolkit

12 → NEXT STEPS

ACTION	OWNER	DEADLINE	PRIORITY

● Low priority ● Medium priority ● High priority

✓ FINAL CHECKLIST BEFORE DEPLOYMENT

✓ Objectives & scope clear · users & stakeholders identified

✓ Data & risks assessed · tests planned · acceptance criteria defined

✓ Oversight in place · approval obtained · decision documented

And next: supervise continuously with **ACF-06**, keep **ACF-07 · Kill Switch** close at hand, and document every decision in **ACF-09 · Decision Register**.



RUNNING EXAMPLE

Lendari documents the dossier for its first agent, "Margaux" (Level 1 customer support), prior to deployment. Lead: **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

PAGE 1/6

1 IDENTITY OF THE AGENTIC PROJECT

AGENT NAME

Margaux

ASSOCIATED PROJECT / SERVICE

Customer support & returns · fashion e-commerce

PROJECT LEAD

Camille Roussel · DDAO

TEAM INVOLVED

L. Fontaine (Customer Support), IT/Ops, P. Dubois (DPO)

DOSSIER CREATION DATE

05/26/2026

DOSSIER VERSION

v1.0

2 PURPOSE & OBJECTIVES OF THE AGENT

PRIMARY PURPOSE

What is the fundamental goal of this agent?

Automate Level 1 customer support to respond quickly to simple requests (tracking, returns, refunds < €100) and free up advisors for complex cases.

OPERATIONAL OBJECTIVES

What must the agent concretely accomplish?

1. Resolve $\geq 70\%$ of Level 1 requests without human intervention
2. First-response time < 2 min
3. Process refunds / credits < €100 (supervised)
4. Maintain CSAT $\geq 4.2 / 5$
5. Cleanly escalate any out-of-scope case

EXAMPLES OF AGENTS

- Intelligent customer support agent
- Commercial analytics agent
- Report generation agent
- Competitive intelligence agent

WATCH POINTS

- ! Objectives too broad or vague
- ! Uncontrolled scope
- ! Sensitive data not identified
- ! Lack of stakeholder buy-in



RUNNING EXAMPLE

Lendari · who uses Margaux and how far she acts: users, stakeholders, and scope.

PAGE 2/6

3 👤 USERS & STAKEHOLDERS

END USERS

Who will interact with the agent?

- Lendari customers (website + app)
- Customer support advisors (supervision & handover)

KEY STAKEHOLDERS

Who must be involved or informed?

- Camille Roussel · DDAO (lead)
- Étienne Mercier · CEO (sponsor)
- Paul Dubois · DPO (compliance)
- Léa Fontaine · Head of Customer Support (business)

4 📄 SCOPE & FEATURES

INCLUDED FEATURES

What the agent will do.

- Order & delivery tracking
- Return management & labels
- Refund / credit < €100
- Product & policy FAQ responses

EXCLUDED FEATURES

What the agent will not do.

- Disputes & refunds > €100
- Fraud detection, legal claims
- Exceptional goodwill gestures
- Modification of an already paid order

SCOPE LIMITS

Refund cap €100 / decision · no access to payment data · EU scope only · mandatory escalation in case of doubt.

✓ BEST PRACTICES

- ✓ Start small, with a minimum viable scope.
- ✓ Document every decision and every limit.
- ✓ Test early and often, with users.

⚠️ WATCH POINTS

- ! Scope too broad at the outset.
- ! Blurred boundary between included and excluded.
- ! Edge cases not settled in advance.



RUNNING EXAMPLE

Lendari · the resources of **Margaux**: data mobilized, access sources, and planned architecture.

PAGE 3/6

5 DATA & RESOURCES

REQUIRED DATA

- Order history
- Delivery status (carrier)
- Product catalog & returns policy

SOURCES & ACCESS

- Order ERP (read API)
- Carrier API (tracking)
- Customer CRM (read) · FAQ database

DATA CONSTRAINTS

GDPR : EU customers, EU instance · minimization · no payment data · logging of accesses & decisions.

6 PLANNED ARCHITECTURE & TOOLS

MAIN COMPONENTS

- Governed LLM + rules engine (caps)
- ERP / CRM / carrier connectors
- Escalation module to human support

TOOLS & TECHNOLOGIES

- Lendari customer support platform (EU instance)
- Sovereign model hosted in the EU
- Internal APIs + kill switch (ACF-07)

INTEGRATIONS & INTERFACES

ERP (orders) · CRM (customers) · Carrier API (tracking) · Slack #sav-incident · supervision dashboard.

EXPECTED DELIVERABLES

- Functional specification of the agent
- Technical documentation
- Test & deployment plan
- Supervision plan + user guide

BEST PRACTICES

- ✓ Prefer a sovereign model & hosting.
- ✓ Give the agent the minimum useful data.
- ✓ Plan the kill switch from the architecture stage.



7 DEPLOYMENT PLAN

STAGE	DESCRIPTION	OWNER	STATUS
Preparation	Scoping & first agent dossier	C. Roussel	● Completed
Development	Connectors, rules, safeguards	IT / Ops	● Completed
Tests & validation	Scenarios + user tests (UAT)	L. Fontaine	● In progress
Pilot deployment	10% of Level 1 customer support traffic	C. Roussel	● To do
Full deployment	100% of Level 1 (if KPIs OK)	C. Roussel	● To do
Monitoring & optimization	Weekly KPI review	L. Fontaine	● To do

8 RISKS & MITIGATIONS

IDENTIFIED RISK	IMPACT	PROB.	MITIGATION PLAN	OWNER
Hallucination in a response	High	3	Sourced responses + escalation in case of doubt	C. Roussel
Erroneous refund	Critical	2	€100 cap + kill switch	C. Roussel
Customer data leak	Critical	2	EU instance + minimization + RLS	P. Dubois
Behavioral drift	High	2	Daily supervision during week 1 + alerts	C. Roussel

Impact: Low / Moderate / High / Critical · Probability: 1 (low) to 5 (high)



RUNNING EXAMPLE

Lendari : test plan, acceptance criteria, and supervision setup for Margaux.

PAGE 5/6

9 TESTS & VALIDATION

PLANNED TEST TYPES

- ✓ Functional & integration
- ✓ Security & performance
- ✓ User acceptance tests (UAT)

ACCEPTANCE CRITERIA

1. **≥ 70% autonomous resolution**
2. **0 refunds > €100 without validation**
3. **First-response time < 2 min (p95)**
4. **Error rate < 2% · escalation 100% out-of-scope**

10 SUPERVISION & IMPROVEMENT PLAN

SUPERVISION

- Real-time dashboard
- Daily during week 1, then weekly

TRACKING METRICS

- Resolution, CSAT, time
- Error, escalation, € refunded

CONTINUOUS IMPROVEMENT

- Weekly review + rule tuning
- FAQ enrichment

✓ BEST PRACTICES

- ✓ Test in real conditions before the pilot.
- ✓ One acceptance criterion = one measurable threshold.
- ✓ Supervise closely during the first week.

⚠ WATCH POINTS

- ! Deploying without a representative test set.
- ! Metrics not instrumented from the outset.
- ! Supervision that fades after launch.



RUNNING EXAMPLE

Lendari : formal validation, deployment decision, and next steps before launching Margaux.

PAGE 6/6

11 VALIDATION & APPROVAL

ROLE	NAME	DATE	SIGN-OFF
Project / AI lead	Camille Roussel · DDAO	06/03/2026	✓
Compliance lead	Paul Dubois · DPO	06/03/2026	✓
Sponsor / decision-maker	Étienne Mercier · CEO	06/03/2026	✓

DEPLOYMENT DECISION

- Approved with reservations - 10% pilot deployment first; full deployment conditional on pilot KPIs.

12 → NEXT STEPS

ACTION	OWNER	DEADLINE	STATUS
Launch the 10% pilot	C. Roussel	06/09/2026	To do
Pilot KPI review (D+7)	L. Fontaine	06/16/2026	To do
Go / No-go full deployment	É. Mercier	06/20/2026	To do

✓ FINAL CHECKLIST BEFORE DEPLOYMENT

- ✓ Objectives & scope clear · users & stakeholders identified
- ✓ Data & risks assessed · tests planned · acceptance criteria defined
- ✓ Supervision in place · approval obtained · decision tracked

What's next: supervise continuously with **ACF-06**, keep **ACF-07** · **Kill Switch** within reach, and log every decision in **ACF-09** · **Decision Register**.



RUNNING EXAMPLE

Rhena documents the dossier for its first agent, "ANNA" (level 1 home claims), prior to deployment. Lead: **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/6

1 IDENTITY OF THE AGENTIC PROJECT

AGENT NAME

ANNA

ASSOCIATED PROJECT / SERVICE

Level 1 home claims · property and casualty

PROJECT LEAD

Dr Katrin Ostermann · DDAO

TEAM INVOLVED

I. Baumgartner (home claims), IT/Ops, S. Lindqvist (DPO)

DOSSIER CREATION DATE

01/14/2027

DOSSIER VERSION

v1.0

2 PURPOSE & OBJECTIVES OF THE AGENT

PRIMARY PURPOSE

What is the fundamental goal of this agent?

Handle level 1 home claims to settle simple files quickly (escape of water, broken glass, simple theft, settlements < €2,500) and free up handlers for complex files.

OPERATIONAL OBJECTIVES

What must the agent concretely accomplish?

1. Settle ≥ 60% of level 1 files without human intervention
2. First-response time < 4h
3. Process settlements < €2,500 (supervised)
4. Maintain policyholder satisfaction ≥ 4.0 / 5
5. Cleanly escalate any out-of-scope file

EXAMPLES OF AGENTS

- Intelligent customer support agent
- Commercial analytics agent
- Report generation agent
- Competitive intelligence agent

WATCH POINTS

- ! Objectives too broad or vague
- ! Uncontrolled scope
- ! Sensitive data not identified
- ! Lack of stakeholder buy-in



3 👤 USERS & STAKEHOLDERS

END USERS

Who will interact with the agent?

- Rhena policyholders (portal + app)
- Claims handlers (supervision & handover)

KEY STAKEHOLDERS

Who must be involved or informed?

- Dr Katrin Ostermann · DDAO (lead)
- Andreas Voll · Management Board (sponsor)
- Sofia Lindqvist · DPO (compliance)
- Ines Baumgartner · Head of Home Claims (business)

4 T SCOPE & FEATURES

INCLUDED FEATURES

What the agent will do.

- File tracking & supporting documents
- Depreciation & excess calculation
- Settlement < €2,500
- Answers on the cover in the policy

EXCLUDED FEATURES

What the agent will not do.

- Refusal of cover & settlements > €2,500
- Fraud, litigation, disputed assessment
- Bodily injury, uninhabitable home
- Modification of a policy in force

SCOPE LIMITS

Settlement cap €2,500 / decision · no bank access · EU scope only · mandatory escalation in case of doubt.

✓ BEST PRACTICES

- ✓ Start small, with a minimum viable scope.
- ✓ Document every decision and every limit.
- ✓ Test early and often, with users.

⚠ WATCH POINTS

- ! Scope too broad at the outset.
- ! Blurred boundary between included and excluded.
- ! Edge cases not settled in advance.



RUNNING EXAMPLE

Rhena · the resources of **ANNA**: data mobilized, access sources, and planned architecture.

PAGE 3/6

5 DATA & RESOURCES

REQUIRED DATA

- Claims history
- Depreciation schedule in force
- Policy, cover & excesses

SOURCES & ACCESS

- Policy repository (read API)
- Claims portal (notifications)
- Versioned schedule (read only)

DATA CONSTRAINTS

GDPR : EU policyholders, Frankfurt region · minimization · no bank data · logging of accesses & decisions.

6 PLANNED ARCHITECTURE & TOOLS

MAIN COMPONENTS

- Governed LLM + rules engine (caps)
- Policy / claims / schedule connectors
- Escalation module to human handlers

TOOLS & TECHNOLOGIES

- Claims portal (Frankfurt region)
- Sovereign model hosted in the EU
- Internal APIs + stop mechanism (ACF-07)

INTEGRATIONS & INTERFACES

Policy repository · claims portal · versioned schedule · incident channel · supervision dashboard.

EXPECTED DELIVERABLES

- Functional specification of the agent
- Technical documentation
- Test & deployment plan
- Supervision plan + user guide

BEST PRACTICES

- ✓ Prefer a sovereign model & hosting.
- ✓ Give the agent the minimum useful data.
- ✓ Plan the kill switch from the architecture stage.



7 DEPLOYMENT PLAN

STAGE	DESCRIPTION	OWNER	STATUS
Preparation	Scoping & first agent dossier	K. Ostermann	● Completed
Development	Connectors, rules, safeguards	IT / Ops	● Completed
Tests & validation	Scenarios + user tests (UAT)	I. Baumgartner	● In progress
Pilot deployment	5% of level 1 files	K. Ostermann	● To do
Full deployment	40% of level 1 (if KPIs OK)	K. Ostermann	● To do
Monitoring & optimization	Weekly KPI review	I. Baumgartner	● To do

8 RISKS & MITIGATIONS

IDENTIFIED RISK	IMPACT	PROB.	MITIGATION PLAN	OWNER
Hallucination in a response	High	3	Sourced responses + escalation in case of doubt	K. Ostermann
Erroneous settlement	Critical	2	€2,500 cap + versioned schedule	K. Ostermann
Policyholder data leak	Critical	2	Frankfurt region + minimization + RLS	S. Lindqvist
Out-of-date depreciation schedule	Critical	3	Version logged on every decision	K. Ostermann

Impact: Low / Moderate / High / Critical · Probability: 1 (low) to 5 (high)



RUNNING EXAMPLE

Rhena : test plan, acceptance criteria, and supervision setup for ANNA.

PAGE 5/6

9 TESTS & VALIDATION

PLANNED TEST TYPES

- ✓ Functional & integration
- ✓ Security & performance
- ✓ User acceptance tests (UAT)

ACCEPTANCE CRITERIA

1. **≥ 60% autonomous settlement**
2. **0 settlements > €2,500 without validation**
3. **First-response time < 4h (p95)**
4. **Error rate < 1.5% · escalation 100% out-of-scope**

10 SUPERVISION & IMPROVEMENT PLAN

SUPERVISION

- Real-time dashboard
- Daily during month 1, then weekly

TRACKING METRICS

- Settlement, satisfaction, time
- Error, escalation, € settled

CONTINUOUS IMPROVEMENT

- Weekly review + rule tuning
- Schedule reconciliation

✓ BEST PRACTICES

- ✓ Test in real conditions before the pilot.
- ✓ One acceptance criterion = one measurable threshold.
- ✓ Supervise closely during the first week.

⚠ WATCH POINTS

- ! Deploying without a representative test set.
- ! Metrics not instrumented from the outset.
- ! Supervision that fades after launch.



RUNNING EXAMPLE

Rhena : formal validation, deployment decision, and next steps before launching ANNA.

PAGE 6/6

11 VALIDATION & APPROVAL

ROLE	NAME	DATE	SIGN-OFF
Project / AI lead	Dr Katrin Ostermann · DDAO	01/14/2027	✓
Compliance lead	Sofia Lindqvist · DPO	01/14/2027	✓
Sponsor / decision-maker	Andreas Voll · Management Board	01/14/2027	✓

DEPLOYMENT DECISION

- Approved with reservations - 5% pilot deployment first; any widening conditional on pilot KPIs.

12 → NEXT STEPS

ACTION	OWNER	DEADLINE	STATUS
Launch the 5% pilot	K. Ostermann	02/01/2027	☐ To do
Pilot KPI review (D+28)	I. Baumgartner	03/01/2027	☐ To do
Go / No-go widening	A. Voll	03/08/2027	☐ To do

✓ FINAL CHECKLIST BEFORE DEPLOYMENT

- ✓ Objectives & scope clear · users & stakeholders identified
- ✓ Data & risks assessed · tests planned · acceptance criteria defined
- ✓ Supervision in place · approval obtained · decision tracked

What's next: supervise continuously with **ACF-06**, keep **ACF-07** · **Kill Switch** within reach, and log every decision in **ACF-09** · **Decision Register**.



LEVEL

DURATION

USE

TRACK

Operational

1 h 45 - 3 h

Traceability / compliance

Toolkit

OBJECTIVE

Ensure the traceability of every significant decision made by an agent - or concerning one. The register documents **decisions, validations, incidents, and follow-up actions** to guarantee transparency, compliance (AI Act art. 12, DORA), and accountability.

1 DECISIONS & INCIDENTS REGISTER

DATE / TIME	AGENT	DECISION / EVENT	RISK	VALIDATED BY	OUTCOME / IMPACT	FOLLOW-UP

Critical = immediate halt **High** = 24 h review **Moderate** = weekly review **Low** = log only **Very low** = sampling



2 NOTES & CONTEXT

3 REGISTER GOVERNANCE RULES



WHY MAINTAIN THIS REGISTER?

- AI Act requirement (art. 12 · logging) & DORA
- Evidence of governance for fiduciary audit
- No record, no accountability



CONFIDENTIALITY

- ! Contains sensitive data (customers, amounts)
- ! Restricted access: DDAO (Delegated Decision Agent Officer), DPO, executive leadership
- ! Retention & purge per GDPR policy



REVIEW FREQUENCY

- ✓ Critical / High: review within 24 h
- ✓ Moderate: weekly review
- ✓ Monthly overall review (DDAO + DPO)

Next: turn tracked incidents into an improvement plan with **ACF-10**, trigger the **ACF-07 · Kill Switch** if a threshold is crossed, and tie each decision back to the **ACF-13 mandate**.



RUNNING EXAMPLE

Lendari logs the decisions of its agent "Margaux" (Tier 1 customer service).
Maintained by **Camille Roussel** (DDAO , Delegated Decision Agent Officer),
supervised together with **Léa Fontaine** (Customer Service).

1 DECISIONS & INCIDENTS REGISTER

DATE / TIME	AGENT	DECISION / EVENT	RISK	VALIDATED BY	OUTCOME / IMPACT	FOLLOW-UP
06/03 09:14	Margaux	Refund €78 · defective item returned	Moderate	Auto (supervised)	Refunded, customer notified < 2 min	Log · weekly review
06/03 11:47	Margaux	Refund request €240 · above the €100 cap	High	Co-decision L. Fontaine	Escalated, approved after invoice check	24-hour review (mandate ACF-13)
06/04 15:30	Margaux	Goodwill gesture -15% on delivery dispute	Moderate	Co-decision L. Fontaine	€32 credit note issued	Satisfaction D+7
06/05 08:02	Margaux	Error rate 2.3% > 2% threshold	Critical	Auto → alert C. Roussel (DDAO)	Kill switch armed, agent suspended 40 min	Post-mortem ACF-10
06/05 10:20	Margaux	Resumed operations after prompt correction	High	C. Roussel (DDAO)	Restart OK, rate dropped to 0.7%	24-hour review
06/05 16:05	Margaux	Customer dispute threatening a negative public review	High	Co-decision L. Fontaine	Out of scope → handled by a human advisor	Closed D+1

Critical

 = immediate halt

High

 = 24-hour review

Moderate

 = weekly review

Low

 = log only

Very low

 = sampling



RUNNING EXAMPLE

Weekly context notes and governance rules for the Lendari register.

PAGE 2/2

2 NOTES & CONTEXT

Incident #4 (06/05). Error spike caused by a faulty catalog synchronization : false "item unavailable" responses. Fix: resynchronization + safeguard added to the prompt. No customer loss. Full post-mortem → **ACF-10**.

Mandate compliance. All refunds > €100 this week were escalated in line with the mandate (**ACF-13**). No "goodwill gesture" decision taken autonomously without co-decision.

Indicators. Weekly CSAT: 4.4 / 5. Tier 1 resolution rate without human intervention: 73%.

3 REGISTER GOVERNANCE RULES



WHY KEEP THIS REGISTER?

- AI Act requirement (art. 12 · logging) & DORA
- Evidence of governance for fiduciary audit
- No trail, no accountability



CONFIDENTIALITY

- ! Contains sensitive data (customers, amounts)
- ! Restricted access: DDAO, DPO, executive management
- ! Retention & purge per GDPR policy



REVIEW FREQUENCY

- ✓ Critical / High: review within 24 hours
- ✓ Moderate: weekly review
- ✓ Monthly overall review (DDAO + DPO)

Next step: turn tracked incidents into an improvement plan with **ACF-10**, trigger **ACF-07 · Kill Switch** if a threshold is crossed, and tie each decision back to the **ACF-13 mandate**.



RUNNING EXAMPLE

Rhena logs the decisions of its agent "ANNA" (level 1 home claims). Maintained by Dr Katrin Ostermann (DDAO , Delegated Decision Agent Officer), supervised together with Ines Baumgartner (Customer Service).

PAGE 1/2

1 DECISIONS & INCIDENTS REGISTER

DATE / TIME	AGENT	DECISION / EVENT	RISK	VALIDATED BY	OUTCOME / IMPACT	FOLLOW-UP
03/11 08:20	ANNA	Settlement €1,920 · escape of water, 4-year-old contents	● Moderate	Auto (supervised)	Settled, policyholder notified < 4h	Log · weekly review
03/11 11:47	ANNA	Claim estimated €3,400 · above the €2,500 cap	● High	Co-decision I. Baumgartner	Escalated, adjuster appointed	24-hour review (mandate ACF-13)
03/11 14:05	ANNA	Complaint about the depreciation rate applied	● Moderate	Co-decision I. Baumgartner	3 further complaints within 20 min	Immediate analysis
03/11 14:22	ANNA	Depreciation schedule still on the 2026 grid	● Critical	Auto → alert K. Ostermann (DDAO)	Stop: Europe 3 min, subsidiary 43 min	Post- mortem ACF-10
03/15 09:00	ANNA	Restart at 5% after the schedule fix	● High	K. Ostermann (DDAO)	Restart OK, 430 files put right	24-hour review
03/15 16:05	ANNA	Policyholder threatening to go to the Ombudsmann	● High	Co-decision I. Baumgartner	Out of scope → handled by a human handler	Closed D+1

● Critical = immediate halt ● High = 24-hour review ● Moderate = weekly review ● Low = log only
● Very low = sampling



2 NOTES & CONTEXT

Incident #4 (03/11). The 1 March tariff revision was not carried into the depreciation schedule : the contents rate stayed at 12% instead of 10%. Fix: schedule corrected, version logged on every decision. 430 policyholders paid the balance, €77,400. Full post-mortem → **ACF-10**.

Mandate compliance. All settlements > €2,500 this week were escalated in line with the mandate (**ACF-13**). No refusal of cover pronounced autonomously, without co-decision.

Indicators. Weekly policyholder satisfaction: 4.1 / 5. Level 1 settlement rate without human intervention: 54%.

3 REGISTER GOVERNANCE RULES

WHY KEEP THIS REGISTER?

- AI Act requirement (art. 12 · logging) & DORA
- Evidence of governance for fiduciary audit
- No trail, no accountability

CONFIDENTIALITY

- ! Contains sensitive data (customers, amounts)
- ! Restricted access: DDAO, DPO, executive management
- ! Retention & purge per GDPR policy

REVIEW FREQUENCY

- ✓ Critical / High: review within 24 hours
- ✓ Moderate: weekly review
- ✓ Monthly overall review (DDAO + DPO)

Next step: turn tracked incidents into an improvement plan with **ACF-10**, trigger **ACF-07 · Kill Switch** if a threshold is crossed, and tie each decision back to the **ACF-13 mandate**.



OBJECTIVE

Turn an incident, audit or review into a concrete action plan: **identify root causes, define corrective and improvement actions, prioritize them (Effort × Impact)** and organize follow-up through to closure.

1 🔍 CONTEXT & ROOT CAUSE

AGENT IN SCOPE

PERIOD COVERED

TRIGGER (INCIDENT / AUDIT / REVIEW)

DESCRIPTION OF THE TRIGGER

ROOT CAUSE IDENTIFIED

2 🔑 CORRECTIVE ACTIONS IDENTIFIED

ACTION	OWNER	DEADLINE	SUCCESS CRITERION	STATUS

Done

In progress

To do

Blocked · escalation required



3 CONTINUOUS IMPROVEMENT ACTIONS

INITIATIVE	EXPECTED BENEFIT	OWNER	DEADLINE

4 PRIORITIZATION · EFFORT × IMPACT MATRIX

IMPACT →

QUICK WINS
High impact · Low effort

IF TIME PERMITS
Low impact · Low effort

MAJOR PROJECTS
High impact · High effort

AVOID
Low impact · High effort

EFFORT →

TOP 3 PRIORITIES

High impact, reasonable effort.

1

2

3



5 FOLLOW-UP & EVALUATION

REVIEW FREQUENCY

NEXT REVIEW SCHEDULED

FOLLOW-UP OWNER

TRACKING INDICATORS

6 VALIDATION

ACTION PLAN OWNER

DDAO (DELEGATED DECISION
AGENT OFFICER)

DATE

✓ BEST PRACTICES

- ✓ SMART actions: a single owner, a clear deadline
- ✓ Measurable success criterion for every action
- ✓ Address the root cause, not just the symptom
- ✓ Regular reviews until complete closure

⚠ WATCH POINTS

- ! An action without deadline or owner is a dead action
- ! Too many simultaneous actions: prioritize
- ! Fixing the symptom without the root cause
- ! No follow-up means no real improvement

Next: log closed actions in the **ACF-09** register, adjust supervision with **ACF-06**, and revise the **ACF-13 mandate** if thresholds or scope have changed.



RUNNING EXAMPLE

Following **incident #4** (error spike on 06/05, kill switch triggered), **Lendari** is building the action plan for its agent "**Margaux**". Led by **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

1 CONTEXT & ROOT CAUSE

AGENT CONCERNED

Margaux · Tier 1 customer service

PERIOD COVERED

Week of 06/03 to 06/05/2026

TRIGGER

Incident #4 · error spike

DESCRIPTION OF THE TRIGGER

Error rate 2.3% > 2% threshold on 06/05 at 08:02. Kill switch armed (ACF-07), agent suspended for 40 min. Logged in the ACF-09 register.

ROOT CAUSE IDENTIFIED

Faulty catalog feed synchronization → false "item unavailable" responses returned to customers. No data freshness safeguard in place.

2 CORRECTIVE ACTIONS IDENTIFIED

ACTION	OWNER	DUE DATE	SUCCESS CRITERION	STATUS
Resync the catalog feed + freshness monitoring	IT / Ops	06/06	0 desync > 15 min	● Done
Prompt safeguard: check real-time availability before reply	C. Roussel	06/08	Error rate < 1%	● In progress
Lower kill switch alert threshold 2% → 1.5%	C. Roussel	06/09	Alert tested OK	● In progress
Replay 30 "false unavailable" cases in validation	L. Fontaine	06/10	100% replayed	● To do
Contact impacted customers (~10% gesture)	L. Fontaine	06/07	100% contacted	● Done

● Done ● In progress ● To do ● Blocked · escalation required




RUNNING EXAMPLE

Beyond the immediate fix, **Lendari** capitalizes: continuous improvement initiatives and **Effort × Impact** prioritization.

3 CONTINUOUS IMPROVEMENT ACTIONS

INITIATIVE	EXPECTED BENEFIT	OWNER	DUE DATE
Real-time dashboard: error rate + CSAT	Continuous visibility	IT / Ops	06/15
Systematize weekly review of the ACF-09 register	Early detection	C. Roussel	Recurring
Extend freshness safeguards to all feeds	Overall robustness	IT / Ops	06/30
Draft a reusable post-incident playbook	Greater responsiveness	C. Roussel	06/20

4 PRIORITIZATION · EFFORT × IMPACT MATRIX



TOP 3 PRIORITIES

High impact, reasonable effort.

- 1 Real-time prompt safeguard · availability check before reply
- 2 Lower kill switch threshold 2% → 1.5%
- 3 Real-time dashboard for error rate + CSAT



RUNNING EXAMPLE

A plan is only as good as its follow-up: indicators, review frequency and validation by the **DDAO**.

PAGE 3/3

5 FOLLOW-UP & EVALUATION

REVIEW FREQUENCY

Daily (W1) then weekly

NEXT REVIEW SCHEDULED

06/12/2026

FOLLOW-UP OWNER

Camille Roussel (DDAO)

FOLLOW-UP INDICATORS

Error rate < 1% · CSAT ≥ 4.2 / 5 · 0 catalog desync > 15 min · % of corrective actions completed (target 100% by 06/12).

6 VALIDATION

ACTION PLAN OWNER

Léa Fontaine · Customer Service
Lead

DDAO

Camille Roussel

DATE

06/06/2026

✓ BEST PRACTICES

- ✓ SMART actions: one single owner, one due date
- ✓ Measurable success criterion for each action
- ✓ Address the root cause, not only the symptom
- ✓ Regular review until full closure

⚠ WATCH-OUTS

- ! Action with no due date or owner = dead action
- ! Too many simultaneous actions: prioritize
- ! Fixing the symptom without the root cause
- ! No follow-up = no real improvement

Next steps: log closed actions to the **ACF-09** register, adjust supervision with **ACF-06**, and revise the **ACF-13** **mandate** if thresholds or scope have changed.



RUNNING EXAMPLE

Following **incident #4** (out-of-date schedule on 03/11, stop triggered), **Rhena** is building the action plan for its agent **"ANNA"**. Led by **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/3

1 CONTEXT & ROOT CAUSE

AGENT CONCERNED

ANNA · level 1 home claims

PERIOD COVERED

Week of 03/09 to 03/13/2027

TRIGGER

Incident #4 · out-of-date schedule

DESCRIPTION OF THE TRIGGER

Under-settlement detected on 03/11 at 14:05. Stop triggered at 14:22 (ACF-07): Europe in 3 min, subsidiary in 43 min. Logged in the ACF-09 register.

ROOT CAUSE IDENTIFIED

The 1 March tariff revision was not carried into the depreciation schedule. No schedule version logged, no reconciliation with the tariff chain.

2 CORRECTIVE ACTIONS IDENTIFIED

ACTION	OWNER	DUE DATE	SUCCESS CRITERION	STATUS
Fix the schedule and replay the 430 files	IT / Ops	03/13	430 files put right	Done
Log the schedule version on every decision	B. Kessler	03/31	100% of decisions versioned	In progress
Extend stop authority to the US subsidiary	K. Ostermann	04/30	Cut-off < 5 min on both sides	In progress
Tie the schedule to the yearly tariff revision	B. Kessler	04/30	Alert at every revision	To do
Pay the balance to the 430 policyholders	I. Baumgartner	03/20	100% paid, €77,400	Done

Done In progress To do Blocked · escalation required



RUNNING EXAMPLE

Beyond the immediate fix, Rhena capitalizes: continuous improvement initiatives and Effort × Impact prioritization.

PAGE 2/3


3

CONTINUOUS IMPROVEMENT ACTIONS

INITIATIVE	EXPECTED BENEFIT	OWNER	DUE DATE
Dashboard: error rate + schedule drift	Continuous visibility	IT / Ops	04/15
Systematize weekly review of the ACF-09 register	Early detection	K. Ostermann	Recurring
Extend versioning to every business reference	Overall robustness	IT / Ops	06/30
Draft a reusable post-incident playbook	Greater responsiveness	K. Ostermann	04/20

4

PRIORITIZATION · EFFORT × IMPACT MATRIX





RUNNING EXAMPLE

A plan is only as good as its follow-up: indicators, review frequency and validation by the **DDAO**.

PAGE 3/3

5 FOLLOW-UP & EVALUATION

REVIEW FREQUENCY

Daily (W1) then weekly

NEXT REVIEW SCHEDULED

03/20/2027

FOLLOW-UP OWNER

Dr Katrin Ostermann (DDAO)

FOLLOW-UP INDICATORS

Error rate < 1% · policyholder satisfaction ≥ 4.0 / 5 · 0 schedule drift · % of corrective actions completed (target 100% by 04/30).

6 VALIDATION

ACTION PLAN OWNER

Ines Baumgartner · Home Claims
Lead

DDAO

Dr Katrin Ostermann

DATE

03/13/2027

✓ BEST PRACTICES

- ✓ SMART actions: one single owner, one due date
- ✓ Measurable success criterion for each action
- ✓ Address the root cause, not only the symptom
- ✓ Regular review until full closure

⚠ WATCH-OUTS

- ! Action with no due date or owner = dead action
- ! Too many simultaneous actions: prioritize
- ! Fixing the symptom without the root cause
- ! No follow-up = no real improvement

Next steps: log closed actions to the **ACF-09** register, adjust supervision with **ACF-06**, and revise the **ACF-13** **mandate** if thresholds or scope have changed.



OBJECTIVE

Measure the maturity of your agentic governance across **8 domains**, prioritize actions and execute them over **30 days**. **Organizational scale 1→5 (CMMI)** · distinct from the **N0→N3** maturity of an individual agent (ACF-05).

1 GENERAL INFORMATION

ORGANIZATION

AUDIT LEAD

AUDIT DATE

SCOPE ASSESSED

AUDIT OBJECTIVE

2 AUDIT GRID · MATURITY ACROSS 8 DOMAINS

DOMAIN ASSESSED	KEY CRITERIA	MATURITY (1 TO 5)	EVIDENCE / COMMENT
Governance	Defined framework · clear roles & responsibilities	1 2 3 4 5	
Strategy & alignment	Strategic alignment · objectives per use case	1 2 3 4 5	
Data	Quality & availability · classification · GDPR	1 2 3 4 5	
Technology	Fit-for-purpose architecture · mastered tooling · resilience	1 2 3 4 5	
Security & compliance	Security controls · AI Act · risk management	1 2 3 4 5	
Oversight & control	Production monitoring · alerts · kill switch	1 2 3 4 5	
Skills & culture	Internal skills · training · responsible culture	1 2 3 4 5	
Performance & value	Defined KPIs · measured value · continuous improvement	1 2 3 4 5	

- 1

Initial · ad hoc
- 2

Reactive · loosely formalized
- 3

Defined · documented
- 4

Managed · measured & controlled
- 5

Optimal · optimized & innovative

Overall maturity level (average across the 8 domains): / 5



3 SCORE INTERPRETATION

AVERAGE SCORE	LEVEL	INTERPRETATION	RECOMMENDED ACTIONS
1.0 – 1.9	INITIAL	Governance is non-existent or very weak. High risk, insufficient control.	Lay the foundations: framework, roles, baseline policies, awareness.
2.0 – 2.9	REACTIVE	Effort underway but unstructured. Significant vulnerabilities.	Formalize key processes, document and strengthen controls.
3.0 – 3.9	DEFINED	Governance in place but improvable. Consistency to be enhanced.	Strengthen measurement, oversight and strategic alignment.
4.0 – 4.4	MANAGED	Robust and measured governance. Sound risk control.	Optimize effectiveness and value. Automate selected controls.
4.5 – 5.0	OPTIMAL	Benchmark governance. Continuous improvement culture.	Innovate, share best practices, pursue excellence.

4 RESULTS ANALYSIS

✓ STRENGTHS IDENTIFIED	⚠ WEAKNESSES / MAJOR GAPS

5 ACTION PRIORITIZATION

PRIORITY	DOMAIN CONCERNED	KEY ACTION	EFFORT (1-5)	IMPACT (1-5)
1				
2				
3				
4				
5				

EFFORT 1 a few hours, no budget · 2 a few days, in-house · 3 a few weeks, limited budget · 4 several months / external vendor · 5 multi-quarter program

IMPACT 1 cosmetic · 2 limited gain · 3 real gain on one key domain · 4 structural gain across multiple domains · 5 unlocks compliance / sovereignty

→ Prioritize High impact × Low effort (quick wins), then high-impact programs.



6

30-DAY ACTION PLAN

WEEK	ACTIONS TO PERFORM	EXPECTED DELIVERABLE	OWNER
W1 · Analysis & scoping			
W2 · Design			
W3 · Implementation			
W4 · Control & adjustment			

7

SELF-REFLECTION QUESTIONS

What is the principal strength of our agentic governance?

What is the greatest uncontrolled risk?

Which decisions must be taken as a priority?

8

GENERAL RECOMMENDATIONS

- ☐ Define a clear, approved governance framework

☐ Maintain a register of agents and their use cases

☐ Strengthen security and compliance (DPO, DPIA)

☐ Train and raise awareness across teams
- ☐ Appoint a governance lead (DDAO · Delegated Decision Agent Officer)

☐ Standardize data management and classification

☐ Deploy oversight and alerting tools

☐ Measure performance and risk on a regular basis



9 EXECUTIVE SUMMARY

MAIN CONCLUSIONS

TOP PRIORITIES (1 TO 5)

1	
2	
3	
4	
5	

EXPECTED BENEFITS

IMMEDIATE NEXT STEPS

10 AUDIT VALIDATION

AUDITOR

FUNCTION

DATE

SIGN-OFF

11 NEXT AUDIT

RECOMMENDED DATE

RECOMMENDED FREQUENCY

Quarterly

Semi-annual

Annual

Ad hoc

✓ BEST PRACTICES

- ✓ Involve key stakeholders in the audit
- ✓ Rely on evidence, not impressions
- ✓ Assess each domain independently
- ✓ Prioritize high-impact actions

⚠ WATCH-OUTS

- ! Neglecting data quality
- ! Underestimating security and compliance
- ! Lacking visibility on agents in production
- ! An audit without a concrete action plan

What's next: roll out the actions in an **ACF-10 plan**, update the **ACF-13 Agent Mandate**, log decisions in the **ACF-09 register**, and re-measure sovereignty with **ACF-01**.



RUNNING EXAMPLE

One month after deploying "Margaux" and **incident #4**, **Lendari** measures the maturity of its agentic governance. Audit conducted by **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

1 GENERAL INFORMATION

ORGANIZATION

Lendari · fashion e-commerce
(€120M, 90 employees)

AUDIT LEAD

Camille Roussel (DDAO)

AUDIT DATE

06/10/2026

SCOPE ASSESSED

Agent "Margaux" (Tier 1 customer service & returns)
+ Lendari governance framework

AUDIT OBJECTIVE

Measure post-incident #4 maturity and frame a 30-day plan of priority actions

2 AUDIT GRID · MATURITY ACROSS 8 DOMAINS

DOMAIN ASSESSED	KEY CRITERIA	MATURITY (1 TO 5)	EVIDENCE / COMMENT
Governance	Framework defined · clear roles	1 2 3 4 5	Constitution v1.0 · DDAO · committee
Strategy & alignment	Alignment · objectives by use case	1 2 3 4 5	ACF-13 Mandate · KPIs (ACF-06)
Data	Quality · classification · GDPR	1 2 3 4 5	Incident #4: desync · no classification
Technology	Architecture · tools · resilience	1 2 3 4 5	Kill switch ACF-07 · LLM dependency
Security & compliance	Security · AI Act · risk management	1 2 3 4 5	DPO · GDPR up to date · DPIA in progress
Supervision & control	Monitoring · alerts · kill switch	1 2 3 4 5	ACF-09 register · weekly review · alerts
Skills & culture	Skills · training · culture	1 2 3 4 5	Know-how held by 2 people · no training
Performance & value	Metrics · value · improvement	1 2 3 4 5	CSAT 4.2 · -30% tickets · ROI pending

1 Initial · ad hoc 2 Reactive · loosely formalized 3 Defined · documented 4 Managed · measured & controlled
5 Optimal · optimized & innovative

Overall maturity level (average of 8 domains): 2.9 / 5 REACTIVE



RUNNING EXAMPLE

Lendari reads its score (2.9 / 5 · Reactive), identifies its strengths and major gaps, then prioritizes actions.

PAGE 2/5

3

SCORE INTERPRETATION

AVERAGE SCORE	LEVEL	INTERPRETATION	RECOMMENDED ACTIONS
1.0 – 1.9	INITIAL	Governance nonexistent or very weak. High risks, insufficient control.	Lay the foundations: framework, roles, basic policies, awareness.
2.0 – 2.9	REACTIVE	Approach underway but unstructured. Significant vulnerabilities.	Formalize key processes, document and strengthen controls.
3.0 – 3.9	DEFINED	Governance in place but improvable. Consistency to be enhanced.	Strengthen measurement, supervision, and strategic alignment.
4.0 – 4.4	MANAGED	Robust and measured governance. Strong risk control.	Optimize effectiveness and value. Automate certain controls.
4.5 – 5.0	OPTIMAL	Benchmark governance. Continuous improvement culture.	Innovate, share best practices, aim for excellence.

4

RESULTS ANALYSIS

✓ IDENTIFIED STRENGTHS

- ✓ Mature supervision & control (4/5): kill switch + register + review
- ✓ Governance framework in place: Agentic Constitution + DDAO
- ✓ Proven incident response · incident #4 contained in 40 min

⚠ WEAKNESSES / MAJOR GAPS

- ! Immature data governance (2/5) · root cause of incident #4
- ! Skills concentrated in 2 people: business continuity risk
- ! AI Act compliance to be formalized (DPIA still in progress)



RUNNING EXAMPLE

Gaps become a **concrete 30-day action plan**, fueled by self-reflection and governance recommendations.

5 ACTION PRIORITIZATION

PRIORITY	DOMAIN CONCERNED	KEY ACTION	EFFORT (1-5)	IMPACT (1-5)
1	Data	Classification + freshness monitoring of catalog data flows	3	5
2	Skills & culture	Train a DDAO pair + upskill the customer service team	2	4
3	Security & compliance	Finalize the AI Act DPIA for agent Margaux	2	4

EFFORT 1 a few hours, no budget · 2 a few days, in-house · 3 a few weeks, limited budget · 4 several months / external provider · 5 multi-quarter initiative

IMPACT 1 cosmetic · 2 limited gain · 3 real gain in one key domain · 4 structural gain across multiple domains · 5 unlocks compliance / sovereignty

→ Prioritize High Impact × Low Effort (quick wins), then high-impact initiatives.

6 30-DAY ACTION PLAN

WEEK	ACTIONS TO COMPLETE	EXPECTED DELIVERABLE	OWNER
W1 · Analysis & scoping	Map data flows; appoint the DDAO pair	Data flow mapping + roles	C. Roussel
W2 · Design	Design freshness monitoring + training plan	Monitoring specs + training plan	IT / Ops
W3 · Implementation	Deploy monitoring; 1st team training session	Monitoring in production · team trained	IT / Ops · C. Roussel
W4 · Control & adjustment	Re-measure maturity (flash audit) and adjust	Flash audit · target maturity 3.3	C. Roussel



RUNNING EXAMPLE

The audit closes with an **executive summary**, its **validation** by the DDAO, and the scheduling of the **next audit**.

PAGE 4/5

7 ? SELF-REFLECTION QUESTIONS

What is the primary strength of our agentic governance?

Supervision discipline: kill switch armed, register maintained, weekly review.

What is the greatest uncontrolled risk?

Catalog data freshness · direct cause of incident #4.

Which decisions must be made first?

Industrialize data monitoring and expand the DDAO pair to reduce individual dependency.

8 ✓ GENERAL RECOMMENDATIONS

- | | |
|---|---|
| ☒ Define a clear, approved governance framework | ☒ Appoint a governance lead (DDAO) |
| ☒ Maintain a register of agents and their use cases | ☐ Standardize data management and classification |
| ☒ Strengthen security and compliance (DPO, DPIA) | ☒ Deploy supervision and alerting tools |
| ☐ Train and raise team awareness | ☒ Measure performance and risks regularly |

9 📄 EXECUTIVE SUMMARY

KEY FINDINGS

Governance is Reactive (2.9/5): supervision is already mature, but data governance and skills must be strengthened. Incident #4 validated control reflexes and exposed the "data" blind spot.

TOP PRIORITIES (1 TO 5)

- 1 Data freshness monitoring
- 2 DDAO pair + team training
- 3 AI Act DPIA for agent Margaux
- 4 Real-time dashboard
- 5 Standardized data classification

EXPECTED BENEFITS

Reduce data incident risk, secure business continuity (less individual dependency), aim for a "Defined" maturity (3.3) at the next audit.

IMMEDIATE NEXT STEPS

Launch the 30-day plan starting 06/11; schedule the flash audit for 07/10/2026.



30-DAY GOVERNANCE AUDIT

*Assess the maturity of your agentic governance
and build a 30-day action plan*

LEVEL	Operational
DURATION	1 - 1.5 days
USE	Self-diagnostic / audit
TRACK	Toolkit

RUNNING EXAMPLE

The audit closes with an **executive summary**, its **validation** by the DDAO, and the scheduling of the **next audit**.

PAGE 5/5

10 AUDIT VALIDATION

AUDITOR

Camille Roussel

ROLE

DDAO · Director of
Digital Transformation

DATE

06/10/2026

SIGN-OFF

C. R. ✓

11 NEXT AUDIT

RECOMMENDED DATE

07/10/2026 (flash audit) then 09/10/2026 (full)

RECOMMENDED FREQUENCY

Quarterly

Semi-annual

Annual

Ad hoc

✓ BEST PRACTICES

- ✓ Involve key stakeholders in the audit
- ✓ Rely on evidence, not impressions
- ✓ Assess each domain independently
- ✓ Prioritize high-impact actions

⚠ WATCH POINTS

- ! Neglecting data quality
- ! Underestimating security and compliance
- ! Lacking visibility into agents in production
- ! An audit without a concrete action plan

What's next: translate actions into an **ACF-10 plan**, update the **ACF-13 Mandate**, log decisions in the **ACF-09 register**, and re-measure sovereignty with **ACF-01**.



RUNNING EXAMPLE

Six months after deploying "ANNA" and **incident #4, Rhena** measures the maturity of its agentic governance. Audit conducted by **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

1 GENERAL INFORMATION

ORGANIZATION

Rhena · insurance (€75bn)

AUDIT LEAD

Dr Katrin Ostermann (DDAO)

AUDIT DATE

09/30/2027

SCOPE ASSESSED

Agent "ANNA" (level 1 home claims) + Rhena governance framework

AUDIT OBJECTIVE

Measure post-incident #4 maturity and frame a 30-day plan of priority actions

2 AUDIT GRID · MATURITY ACROSS 8 DOMAINS

DOMAIN ASSESSED	KEY CRITERIA	MATURITY (1 TO 5)	EVIDENCE / COMMENT
Governance	Framework defined · clear roles	1 2 3 4 5	Constitution v1.0 · DDAO · committee
Strategy & alignment	Alignment · objectives by use case	1 2 3 4 5	ACF-13 Mandate · KPIs (ACF-06)
Data	Quality · classification · GDPR	1 2 3 4 5	Schedule versioned since 03/31
Technology	Architecture · tools · resilience	1 2 3 4 5	Stop ACF-07 · two platforms, one model
Security & compliance	Security · AI Act · risk management	1 2 3 4 5	DPO · GDPR up to date · DPIA in progress
Supervision & control	Monitoring · alerts · kill switch	1 2 3 4 5	Escalation at 41% · capacity not recomputed
Skills & culture	Skills · training · culture	1 2 3 4 5	Know-how held by 2 people · no training
Performance & value	Metrics · value · improvement	1 2 3 4 5	Satisfaction 4.1 · error 0.9% · time 2h10

1 Initial · ad hoc 2 Reactive · loosely formalized 3 Defined · documented 4 Managed · measured & controlled
5 Optimal · optimized & innovative

Overall maturity level (average of 8 domains): 3.1 / 5 DEFINED



RUNNING EXAMPLE

Rhena reads its score (3.1 / 5 · Defined), identifies its strengths and major gaps, then prioritizes actions.

PAGE 2/5

3 SCORE INTERPRETATION

AVERAGE SCORE	LEVEL	INTERPRETATION	RECOMMENDED ACTIONS
1.0 – 1.9	INITIAL	Governance nonexistent or very weak. High risks, insufficient control.	Lay the foundations: framework, roles, basic policies, awareness.
2.0 – 2.9	REACTIVE	Approach underway but unstructured. Significant vulnerabilities.	Formalize key processes, document and strengthen controls.
3.0 – 3.9	DEFINED	Governance in place but improvable. Consistency to be enhanced.	Strengthen measurement, supervision, and strategic alignment.
4.0 – 4.4	MANAGED	Robust and measured governance. Strong risk control.	Optimize effectiveness and value. Automate certain controls.
4.5 – 5.0	OPTIMAL	Benchmark governance. Continuous improvement culture.	Innovate, share best practices, aim for excellence.

4 RESULTS ANALYSIS

✓ IDENTIFIED STRENGTHS

- ✓ Mature governance (4/5): DDAO, signed mandate, register maintained
- ✓ All quality indicators met (4/5): error, time, satisfaction
- ✓ Proven incident response · Europe cut off in 3 min

⚠ WEAKNESSES / MAJOR GAPS

- ! Supervision overrun (3/5) · escalation at 41% against a 35% target
- ! Skills concentrated in 2 people: business continuity risk
- ! Bias by postcode not established, therefore not ruled out (ACF-P15)



RUNNING EXAMPLE

Gaps become a **concrete 30-day action plan**, fueled by self-reflection and governance recommendations.

PAGE 3/5

5 ACTION PRIORITIZATION

PRIORITY	DOMAIN CONCERNED	KEY ACTION	EFFORT (1-5)	IMPACT (1-5)
1	Supervision	Recompute rework capacity and tighten the perimeter	3	5
2	Skills & culture	Train a DDAO pair + upskill the home claims unit	2	4
3	Security & compliance	Establish the absence of bias by postcode (ACF-P15)	2	4

EFFORT 1 a few hours, no budget · 2 a few days, in-house · 3 a few weeks, limited budget · 4 several months / external provider · 5 multi-quarter initiative

IMPACT 1 cosmetic · 2 limited gain · 3 real gain in one key domain · 4 structural gain across multiple domains · 5 unlocks compliance / sovereignty

→ Prioritize High Impact × Low Effort (quick wins), then high-impact initiatives.

6 30-DAY ACTION PLAN

WEEK	ACTIONS TO COMPLETE	EXPECTED DELIVERABLE	OWNER
W1 · Analysis & scoping	Measure real rework load; appoint the DDAO pair	Measured load + roles	K. Ostermann
W2 · Design	Tighten the delegated perimeter + training plan	Revised perimeter + training plan	IT / Ops
W3 · Implementation	Apply the new perimeter; 1st training session	Perimeter in production · team trained	IT / Ops · K. Ostermann
W4 · Control & adjustment	Re-measure maturity (flash audit) and adjust	Flash audit · target maturity 3.5	K. Ostermann



RUNNING EXAMPLE

The audit closes with an **executive summary**, its **validation** by the DDAO, and the scheduling of the **next audit**.

PAGE 4/5

7 ? SELF-REFLECTION QUESTIONS

What is the primary strength of our agentic governance?

The written framework: DDAO appointed, mandate signed, register maintained, stop tested.

What is the greatest uncontrolled risk?

Supervision load: 672 files reworked a day against a capacity of 492.

Which decisions must be made first?

Tighten the delegated perimeter and expand the DDAO pair to reduce individual dependency.

8 ✓ GENERAL RECOMMENDATIONS

- ☒ Define a clear, approved governance framework
- ☒ Maintain a register of agents and their use cases
- ☒ Strengthen security and compliance (DPO, DPIA)
- ☐ Train and raise team awareness
- ☒ Appoint a governance lead (DDAO)
- ☐ Standardize data management and classification
- ☒ Deploy supervision and alerting tools
- ☒ Measure performance and risks regularly

9 📄 EXECUTIVE SUMMARY

KEY FINDINGS

Governance is Defined (3.1/5): the framework is written and kept, but supervision is overrun and skills must be strengthened. Incident #4 validated control reflexes and exposed the "data" blind spot.

TOP PRIORITIES (1 TO 5)

- 1 Supervision capacity recomputed
- 2 DDAO pair + team training
- 3 Bias study by postcode
- 4 Real-time dashboard
- 5 Delegated perimeter tightened

EXPECTED BENEFITS

Bring escalation below 35%, secure business continuity (less individual dependency), aim for a consolidated "Defined" maturity (3.5) at the next audit.

IMMEDIATE NEXT STEPS

Launch the 30-day plan starting 10/01; schedule the flash audit for 10/30/2027.



30-DAY GOVERNANCE AUDIT

*Assess the maturity of your agentic governance
and build a 30-day action plan*

LEVEL	Operational
DURATION	1 - 1.5 days
USE	Self-diagnostic / audit
TRACK	Toolkit

RUNNING EXAMPLE

The audit closes with an **executive summary**, its **validation** by the DDAO, and the scheduling of the **next audit**.

PAGE 5/5

10 AUDIT VALIDATION

AUDITOR

Dr Katrin Ostermann

ROLE

DDAO · Chief
Compliance Officer

DATE

09/30/2027

SIGN-OFF

K. O. ✓

11 NEXT AUDIT

RECOMMENDED DATE

10/30/2027 (flash audit) then 03/30/2028 (full)

RECOMMENDED FREQUENCY

Quarterly

Semi-annual

Annual

Ad hoc

✓ BEST PRACTICES

- ✓ Involve key stakeholders in the audit
- ✓ Rely on evidence, not impressions
- ✓ Assess each domain independently
- ✓ Prioritize high-impact actions

⚠ WATCH POINTS

- ! Neglecting data quality
- ! Underestimating security and compliance
- ! Lacking visibility into agents in production
- ! An audit without a concrete action plan

What's next: translate actions into an **ACF-10 plan**, update the **ACF-13 Mandate**, log decisions in the **ACF-09 register**, and re-measure sovereignty with **ACF-01**.



OBJECTIVE

Identify, assess and prioritize the risks of your AI agents. Each risk is scored **Likelihood × Severity** (score 1→25), then assigned to one of the **8 ACF domains** (same categories as the **ACF-11** audit) in order to build a prioritized treatment plan.

1 GENERAL INFORMATION

ORGANIZATION

ENTITY / DEPARTMENT

ASSESSMENT LEAD

DATE · SCOPE ASSESSED

AGENTS / SYSTEMS COVERED

2 RISK MATRIX

Position each risk according to its **likelihood** (rows, 5→1) and its **severity** (columns, 1→5). The **score** = Likelihood × Severity (1 to 25). The cell color indicates the risk level (see legend).

SCORING SCALE · HOW TO CHOOSE L AND S (1 TO 5)

LVL.	LIKELIHOOD (L) · THAT THE RISK OCCURS	SEVERITY (S) · IMPACT IF THE RISK OCCURS
1	Very low · hardly ever (< 5 %/yr), no precedent.	Minor · negligible, fixed in routine operations.
2	Low · unlikely (~5–20 %), rare across several years.	Moderate · limited disruption, moderate cost/delay, few users affected.
3	Moderate · possible (~20–50 %), a few times a year.	Serious · real disruption, significant loss, minor non-compliance.
4	High · probable (~50–80 %), recurring / near-misses already observed.	Major · significant harm, large-scale user impact, reportable regulatory breach.
5	Very high · near-certain (> 80 %), already occurred / frequent.	Critical · massive harm, AI Act/GDPR sanction, lasting damage.

Indicative scale : adapt thresholds (% , amounts) to your context. The **L × S** score is then read from the matrix.

		SEVERITY (IMPACT) →				
LIK. ↓		1 MINOR	2 MODERATE	3 SERIOUS	4 MAJOR	5 CRITICAL
	5 VERY HIGH	5	10	15	20	25
	4 HIGH	4	8	12	16	20
	3 MODERATE	3	6	9	12	15
	2 LOW	2	4	6	8	10
	1 VERY LOW	1	2	3	4	5

- **Critical** 20–25 · immediate action ● **High** 13–19 · priority action ● **Moderate** 6–12 · monitoring & plan ● **Low** 1–5 · risk under control



LEVEL

DURATION

USE

TRACK

Operational

1 - 2 days

Assessment / prioritization

Toolkit

3

≡ RISK INVENTORY

How to score: 1 Likelihood L (1–5) 2 Severity S (1–5) 3 Score = L × S
4 Level: 1–5 Low · 6–12 Moderate · 13–19 High · 20–25 Critical · definitions of L and S: [scale p. 1](#).

ID	RISK DESCRIPTION (CAUSE / SCENARIO)	CATEGORY	L	S	SCORE	LEVEL	EXISTING CONTROLS

Categories = the 8 ACF domains (same as the ACF-11 audit): Governance · Strategy & alignment · Data · Technology · Security & compliance · Oversight & control · Skills & culture · Performance & value.

4

⌵ ANALYSIS OF MAJOR RISKS

TOP 5 HIGHEST RISKS

1

2

3

4

5

MAIN IDENTIFIED ROOT CAUSES



LEVEL

DURATION

USE

TRACK

Operational

1 - 2 days

Assessment / prioritization

Toolkit

5 TREATMENT & PRIORITIZATION BY LEVEL

RISK LEVEL	PRIORITY	RECOMMENDED TREATMENT	EXAMPLE ACTIONS
CRITICAL 20–25	P1	Immediate treatment · unacceptable risk in its current state.	Suspend or restrict the agent · strengthen controls · narrow scope.
HIGH 13–19	P2	Priority treatment · reduction required.	Mitigation measures · reinforce oversight · short-term action plan.
MODERATE 6–12	P3	Planned treatment · acceptable under control conditions.	Targeted controls · active monitoring · periodic reassessment.
LOW 1–5	P4	Continuous monitoring · risk under control.	Maintain existing controls · regular monitoring · reassessment.

6 OVERALL ACTION PLAN

PRIO.	MAIN RISK	LEVEL	KEY ACTION	DUE DATE	STATUS	OWNER

✓ BEST PRACTICES

- ✓ Address critical and high risks first
- ✓ Stay factual and rely on data
- ✓ Document every decision and assumption

⚠ POINTS OF VIGILANCE

- ! Do not underestimate data and security risks
- ! Include ethical, legal and reputational risks
- ! Do not confuse inherent risk with residual risk



7 MONITORING & REASSESSMENT

REASSESSMENT FREQUENCY

Monthly

Quarterly

Semiannual

Event-driven

NEXT REASSESSMENT SCHEDULED FOR

TRIGGERS & METHODS

- **Triggers:** major change to the agent or its environment · new incident or breach · regulatory evolution (AI Act).
- **Methods:** review of controls · incident analysis · stakeholder consultation · testing & audit.

8 KEY INDICATORS

RISKS IDENTIFIED

CRITICAL + HIGH

UNDER TREATMENT

OPEN RISKS

AVERAGE LEAD TIME

9 SUMMARY & REPORT

KEY CONCLUSIONS

KEY RECOMMENDATIONS (PRIORITIES)

1

2

3

4

5

EXPECTED BENEFITS FROM RISK MANAGEMENT

10 VALIDATION

VALIDATOR

ROLE

DATE

SIGNATURE

What's next: feed the **ACF-10 action plan**, log decisions in the **ACF-09 register**, verify the **ACF-07 kill switch**, update the **ACF-13 mandate** and loop back with the **ACF-11 audit**.



RUNNING EXAMPLE

Two days after the **ACF-11** audit, **Lendari** maps the risks of agent "Margaux" and its data chain. Assessment led by **Camille Roussel** (DDAO · Delegated Decision Agent Officer).

1 GENERAL INFORMATION

ORGANIZATION	ENTITY / DEPARTMENT	ASSESSMENT LEAD
Lendari · fashion e-commerce (€120M, 90 employees)	Digital transformation / Customer service	Camille Roussel (DDAO)
DATE · SCOPE ASSESSED	AGENTS / SYSTEMS CONCERNED	
06/12/2026 : Agent "Margaux" (Tier-1 customer service & returns) + associated data chain	Agent Margaux (N2) · returns/refunds database · e-commerce platform integration	

2 RISK MATRIX

Each risk is positioned by its **probability** (rows, 5→1) and its **severity** (columns, 1→5). The **score** = Probability × Severity (1 to 25). The pins indicate where Lendari's risks fall.

SCORING SCALE · HOW TO CHOOSE P AND S (1 TO 5)

LVL.	PROBABILITY (P) · THAT THE RISK OCCURS	SEVERITY (S) · IMPACT IF THE RISK OCCURS
1	Very low · almost never (< 5 %/year), no precedent.	Minor · negligible, fixed as routine.
2	Low · unlikely (~5–20 %), rare over several years.	Moderate : limited disruption, modest cost/delay, few users.
3	Moderate · possible (~20–50 %), a few times per year.	Serious : real disruption, noticeable loss, minor non-compliance.
4	High : probable (~50–80 %), recurring / near-misses already observed.	Major : significant harm, large numbers of users affected, reportable regulatory breach.
5	Very high : near-certain (> 80 %), already occurred / frequent.	Critical : massive harm, AI Act/GDPR sanction, lasting damage.

Indicative scale : adapt the thresholds (% , amounts) to your context. The **P × S** score is then read in the matrix.

		SEVERITY (IMPACT) →				
PROB. ↓		1 MINOR	2 MODERATE	3 SERIOUS	4 MAJOR	5 CRITICAL
	5 VERY HIGH	5	10	15	20	25
4 HIGH	4	4	8	12	16	20 R1
3 MODERATE	3	3	6	9 R5	12 R3-R4	15
2 LOW	2	2	4	6 R7	8 R2-R6	10
1 VERY LOW	1	1	2	3	4	5



RUNNING EXAMPLE

Lendari inventories **7 risks**, scores them Probability × Severity, then isolates the major risks and their root causes.

PAGE 2/4

3 RISK INVENTORY

How to score: 1 Probability P (1–5) 2 Severity S (1–5) 3 Score = P × S
4 Level: 1–5 Low · 6–12 Moderate · 13–19 High · 20–25 Critical · definitions of P and S: [scale p. 1](#).

ID	RISK DESCRIPTION (CAUSE / SCENARIO)	CATEGORY	P	S	SCORE	LEVEL	EXISTING CONTROLS
R1	Refund data desynchronization	Data	4	5	20	CRITICAL	Root cause of incident #4 · monitoring in progress (ACF-10)
R2	Refund beyond the €100 cap	Supervision & control	2	4	8	MODERATE	Hard cap €100 · weekly review (ACF-09)
R3	Knowledge concentrated in 2 people	Skills & culture	3	4	12		DDAO pairing in progress (ACF-10)
R4	DPIA / AI Act for Margaux not formalized	Security & compliance	3	4	12	MODERATE	DPIA scheduled S3 (ACF-11 plan)
R5	Inappropriate customer service response (tone, hallucination)	Technology	3	3	9		Daily supervision S1 · tone guardrail
R6	Post-hoc incident detection (no real-time)	Supervision & control	2	4	8	MODERATE	Kill switch > 2% error, but delayed detection
R7	Model drift after vendor update	Technology	2	3	6		Non-regression tests before deployment

Categories = the 8 ACF domains (same as the ACF-11 audit).

4 ANALYSIS OF MAJOR RISKS

TOP 5 HIGHEST RISKS

- 1 R1 · Data desynchronization (20 · Critical)
- 2 R3 · Dependency on 2 people (12 · Moderate)
- 3 R4 · DPIA / AI Act not formalized (12 · Moderate)
- 4 R5 · Inappropriate customer service response (9 · Moderate)
- 5 R2 / R6 · Cap & delayed detection (8 · Moderate)

MAIN ROOT CAUSES IDENTIFIED

- Refund data not monitored (cause of incident #4)
- Young governance, concentrated on two people
- AI Act / DPIA compliance not yet formalized
- Delayed incident detection · no real-time alerting



RUNNING EXAMPLE

Each risk is given a **treatment according to its level**, broken down into a **prioritized global action plan**.

PAGE 3/4

5 TREATMENT & PRIORITIZATION BY LEVEL

RISK LEVEL	PRIORITY	RECOMMENDED TREATMENT	EXAMPLE ACTIONS
CRITICAL 20–25	P1	Immediate treatment · risk unacceptable as it stands.	Suspend or restrict the agent · strengthen controls · narrow the scope.
HIGH 13–19	P2	Priority treatment · reduction required.	Mitigation measures · strengthen supervision · short-term action plan.
MODERATE 6–12	P3	Planned treatment · acceptable subject to control conditions.	Targeted controls · active monitoring · periodic re-assessment.
LOW 1–5	P4	Ongoing monitoring · risk under control.	Maintain existing controls · regular monitoring · re-assessment.

6 GLOBAL ACTION PLAN

PRIO.	MAIN RISK	LEVEL	KEY ACTION	DEADLINE	STATUS	OWNER
1	R1 · Refund data desync.	CRIT.	Deploy freshness monitoring + classification	06/30/2026	IN PROGRESS	C. Roussel
2	R3 · Dependency on 2 people	MOD.	Establish DDAO pairing + training plan	07/15/2026	TO DO	C. Roussel
3	R4 · DPIA / AI Act Margaux		Formalize the DPIA + compliance registry	07/15/2026	TO DO	P. Dubois
4	R5 · Inappropriate customer service response	MOD.	Tone guardrail + weekly sample review	06/30/2026	IN PROGRESS	L. Fontaine
5	R6 · Post-hoc detection		Implement real-time alerting on error rate	07/31/2026	TO DO	C. Roussel

✓ BEST PRACTICES

- ✓ Address critical and high risks first
- ✓ Be factual and rely on data
- ✓ Document all decisions and assumptions

⚠ WATCH-OUTS

- ! Do not underestimate data and security risks
- ! Include ethical, legal and reputational risks
- ! Do not confuse inherent risk with residual risk



RUNNING EXAMPLE

The assessment closes with **monitoring**, **key indicators**, a **summary** and its **validation** by the DDAO.

7 MONITORING & RE-ASSESSMENT

RE-ASSESSMENT FREQUENCY

Monthly

Quarterly

Semi-annually

Event-driven


NEXT RE-ASSESSMENT SCHEDULED FOR

09/12/2026 (or earlier if triggered)

TRIGGERS & METHODS

- **Triggers:** major change to the agent or its environment · new incident or vulnerability · regulatory evolution (AI Act).
- **Methods:** review of controls · incident analysis · stakeholder consultation · tests & audit.

8 KEY INDICATORS



9 SUMMARY & REPORT

MAIN CONCLUSIONS

A single critical risk (data desynchronization , cause of incident #4), already under treatment. Six moderate risks under control. Profile consistent with the "Reactive" maturity (ACF-11).

KEY RECOMMENDATIONS (PRIORITIES)

- 1 Data freshness monitoring
- 2 DDAO pairing + training
- 3 DPIA AI Act for Margaux
- 4 Real-time error rate alerting
- 5 Model non-regression tests

EXPECTED BENEFITS OF RISK MANAGEMENT

Reduce residual risk, secure AI Act compliance, strengthen human resilience (less individual dependency) and detect incidents earlier.

10 VALIDATION

VALIDATOR Camille Roussel	ROLE DDAO · Director of digital transformation	DATE 06/12/2026	SIGNATURE C. R. ✓
------------------------------	---	--------------------	----------------------

Next steps: feed the **ACF-10 action plan**, log decisions in the **ACF-09 registry**, verify the **ACF-07 kill switch**, update the **ACF-13 mandate** and loop back to the **ACF-11 audit**.



RUNNING EXAMPLE

Two days after the **ACF-11** audit, **Rhena** maps the risks of agent "ANNA" and its data chain. Assessment led by **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/4

1 GENERAL INFORMATION

ORGANIZATION

Rhena · insurance (€75bn)

ENTITY / DEPARTMENT

Compliance / home claims unit

ASSESSMENT LEAD

Dr Katrin Ostermann (DDAO)

DATE · SCOPE ASSESSED

10/02/2027 : Agent "ANNA" (level 1 home claims) +
associated data chain

AGENTS / SYSTEMS CONCERNED

Agent ANNA (N2) · versioned depreciation schedule ·
policy repository

2 RISK MATRIX

Each risk is positioned by its **probability** (rows, 5→1) and its **severity** (columns, 1→5). The **score** = Probability × Severity (1 to 25). The pins indicate where Rhena's risks fall.

SCORING SCALE · HOW TO CHOOSE P AND S (1 TO 5)

LVL.	PROBABILITY (P) · THAT THE RISK OCCURS	SEVERITY (S) · IMPACT IF THE RISK OCCURS
1	Very low · almost never (< 5 %/year), no precedent.	Minor · negligible, fixed as routine.
2	Low · unlikely (~5–20 %), rare over several years.	Moderate : limited disruption, modest cost/delay, few users.
3	Moderate · possible (~20–50 %), a few times per year.	Serious : real disruption, noticeable loss, minor non-compliance.
4	High : probable (~50–80 %), recurring / near-misses already observed.	Major : significant harm, large numbers of users affected, reportable regulatory breach.
5	Very high : near-certain (> 80 %), already occurred / frequent.	Critical : massive harm, AI Act/GDPR sanction, lasting damage.

Indicative scale : adapt the thresholds (% , amounts) to your context. The **P × S** score is then read in the matrix.

		SEVERITY (IMPACT) →				
		1 MINOR	2 MODERATE	3 SERIOUS	4 MAJOR	5 CRITICAL
PROB. ↓	5 VERY HIGH	5	10	15	20	25
	4 HIGH	4	8	12	16	20 R1
	3 MODERATE	3	6	9 R5	12 R3·R4	15
	2 LOW	2	4	6 R7	8 R2·R6	10
	1 VERY LOW	1	2	3	4	5



RUNNING EXAMPLE

Rhena inventories **7 risks**, scores them Probability × Severity, then isolates the major risks and their root causes.

PAGE 2/4

3 RISK INVENTORY

How to score: 1 Probability P (1–5) 2 Severity S (1–5) 3 Score = P × S

4 Level: 1–5 Low · 6–12 Moderate · 13–19 High · 20–25 Critical · definitions of P and S: [scale p. 1](#).

ID	RISK DESCRIPTION (CAUSE / SCENARIO)	CATEGORY	P	S	SCORE	LEVEL	EXISTING CONTROLS
R1	Schedule out of sync with the tariff revision	Data	4	5	20	CRITICAL	Root cause of incident #4 · monitoring in progress (ACF-10)
R2	Settlement beyond the €2,500 cap	Supervision & control	2	4	8	MODERATE	Hard cap €100 · weekly review (ACF-09)
R3	Knowledge concentrated in 2 people	Skills & culture	3	4	12		DDAO pairing in progress (ACF-10)
R4	Bias by postcode not established	Security & compliance	3	4	12	MODERATE	DPIA scheduled S3 (ACF-11 plan)
R5	Inappropriate response to a policyholder (tone, hallucination)	Technology	3	3	9		Daily supervision S1 · tone guardrail
R6	Stop authority not extended to the subsidiary	Supervision & control	2	4	8	MODERATE	Kill switch > 2% error, but delayed detection
R7	Model drift after vendor update	Technology	2	3	6		Non-regression tests before deployment

Categories = the 8 ACF domains (same as the ACF-11 audit).

4 ANALYSIS OF MAJOR RISKS

TOP 5 HIGHEST RISKS

- 1 R1 · Schedule out of sync (20 · Critical)
- 2 R3 · Dependency on 2 people (12 · Moderate)
- 3 R4 · Bias by postcode not established (12 · Moderate)
- 4 R5 · Inappropriate response to a policyholder (9 · Moderate)
- 5 R2 / R6 · Cap & stop not extended (8 · Moderate)

MAIN ROOT CAUSES IDENTIFIED

- Schedule not tied to the tariff chain (cause of incident #4)
- Young governance, concentrated on two people
- Absence of bias not established, therefore not ruled out
- Stop limited to Europe · the subsidiary is not covered



RUNNING EXAMPLE

Each risk is given a **treatment according to its level**, broken down into a **prioritized global action plan**.

5 TREATMENT & PRIORITIZATION BY LEVEL

RISK LEVEL	PRIORITY	RECOMMENDED TREATMENT	EXAMPLE ACTIONS
CRITICAL 20–25	P1	Immediate treatment · risk unacceptable as it stands.	Suspend or restrict the agent · strengthen controls · narrow the scope.
HIGH 13–19	P2	Priority treatment · reduction required.	Mitigation measures · strengthen supervision · short-term action plan.
MODERATE 6–12	P3	Planned treatment · acceptable subject to control conditions.	Targeted controls · active monitoring · periodic re-assessment.
LOW 1–5	P4	Ongoing monitoring · risk under control.	Maintain existing controls · regular monitoring · re-assessment.

6 GLOBAL ACTION PLAN

PRIO.	MAIN RISK	LEVEL	KEY ACTION	DEADLINE	STATUS	OWNER
1	R1 · Schedule out of sync	CRIT.	Log the schedule version + tariff reconciliation	03/31/2027	IN PROGRESS	K. Ostermann
2	R3 · Dependency on 2 people	MOD.	Establish DDAO pairing + training plan	12/31/2027	TO DO	K. Ostermann
3	R4 · Bias by postcode		Bias study + compliance registry	01/15/2028	TO DO	S. Lindqvist
4	R5 · Inappropriate response	MOD.	Tone guardrail + weekly sample review	10/31/2027	IN PROGRESS	I. Baumgartner
5	R6 · Stop not extended		Extend stop authority to both entities	04/30/2027	TO DO	K. Ostermann

✓ BEST PRACTICES

- ✓ Address critical and high risks first
- ✓ Be factual and rely on data
- ✓ Document all decisions and assumptions

⚠ WATCH-OUTS

- ! Do not underestimate data and security risks
- ! Include ethical, legal and reputational risks
- ! Do not confuse inherent risk with residual risk



RUNNING EXAMPLE

The assessment closes with **monitoring**, **key indicators**, a **summary** and its **validation** by the DDAO.

PAGE 4/4

7 MONITORING & RE-ASSESSMENT

RE-ASSESSMENT FREQUENCY

Monthly

Quarterly

Semi-annually

Event-driven


NEXT RE-ASSESSMENT SCHEDULED FOR

01/02/2028 (or earlier if triggered)

TRIGGERS & METHODS

- **Triggers:** major change to the agent or its environment · new incident or vulnerability · regulatory evolution (AI Act).
- **Methods:** review of controls · incident analysis · stakeholder consultation · tests & audit.

8 KEY INDICATORS



9 SUMMARY & REPORT

MAIN CONCLUSIONS

A single critical risk (schedule out of sync · cause of incident #4), already treated. Six moderate risks under control. Profile consistent with the "Defined" maturity (ACF-11).

KEY RECOMMENDATIONS (PRIORITIES)

- 1 Schedule version log
- 2 DDAO pairing + training
- 3 Bias study by postcode
- 4 Stop authority across both entities
- 5 Model non-regression tests

EXPECTED BENEFITS OF RISK MANAGEMENT

Reduce residual risk, establish the absence of bias, strengthen human resilience (less individual dependency) and cut off everywhere at once.

10 VALIDATION

VALIDATOR Dr Katrin Ostermann	ROLE DDAO · Chief Compliance Officer	DATE 10/02/2027	SIGNATURE K. O. ✓
----------------------------------	--	--------------------	----------------------

Next steps: feed the **ACF-10 action plan**, log decisions in the **ACF-09 registry**, verify the **ACF-07 kill switch**, update the **ACF-13 mandate** and loop back to the **ACF-11 audit**.



OBJECTIVE

Formalize the **mandate** of an agent: its mission, its authorization scope, its limits, its level of **delegated autonomy**, its data, its supervision and its **official validation**. The mandate is the **enforceable** reference document binding the agent and those who govern it. (DDAO = Delegated Decision Agent Officer.)

1 GENERAL INFORMATION

AGENT NAME

BUSINESS SCOPE / BU

MATURITY (ACF-05)

EFFECTIVE DATE

SCHEDULED REVIEW DATE

MANDATE OWNER (DDAO)

2 AGENT MISSION

PRIMARY MISSION · WHY IT EXISTS

PRIORITY OBJECTIVES (TOP 3)

1

2

3

3 AUTHORIZATION SCOPE

✓ AUTHORIZED ACTIONS

☐☐☐☐☐

⊘ PROHIBITED ACTIONS · RED ZONE

☐☐☐☐☐



LEVEL	Operational
DURATION	1 - 1.5 days
USE	Framing / authorization
TRACK	Toolkit

4 ⊕ LIMITS & CONSTRAINTS

FINANCIAL CAP

AUTHORIZED TIME WINDOWS

GEOGRAPHIC / JURISDICTIONAL SCOPE

EXCLUDED CUSTOMER CATEGORIES

5 ⚡ DELEGATED AUTONOMY LEVEL (PER DECISION)

DECISION TYPE	DELEGATED AUTONOMY	THRESHOLD / CONDITION

SUGGESTIVE

 the agent proposes, the human decides

CO-DECISION

 human validation before execution

SUPERVISED AUTONOMOUS

 acts alone, post-hoc control

AUTONOMOUS

 acts alone, without systematic control

Carry over the agent's overall maturity (N0 → N3, ACF-05). Autonomy is set **decision by decision** (ACF-02 / ACF-03) · distinct from maturity.



LEVEL	Operational
DURATION	1 - 1.5 days
USE	Framing / authorization
TRACK	Toolkit

6 DATA & ACCESS

ACCESS LEVEL	DATA COVERED
Read	
Write	
Prohibited	

7 CONTROLS & SUPERVISION

EFFECTIVE HUMAN SUPERVISION

DECISION-QUALITY KPIS

ESCALATION PROCEDURE TO THE DDAO

8 STAKEHOLDERS & AT-RISK GROUPS

PARTIES TO NOTIFY IN CASE OF INCIDENT

AT-RISK USER CATEGORIES

9 TRACEABILITY & REPORTING · LINK ACF-09

LOGGING FORMAT

REPORTING FREQUENCY

RECIPIENTS



10 OFFICIAL VALIDATION - 4 SIGNATURES REQUIRED

BUSINESS OWNER

NAME & SIGNATURE

TECHNICAL OWNER

NAME & SIGNATURE

COMPLIANCE OFFICER / DPO

NAME & SIGNATURE

DDAO

NAME & SIGNATURE

Mandate active · validation date: ----- · next review: ----- . Any change of scope requires a new validation.

✓ BEST PRACTICES

- ✓ List prohibitions as explicitly as authorizations
- ✓ Set autonomy decision by decision, not in bulk
- ✓ Review the mandate on a fixed schedule and after any incident

⚠ WATCH-OUTS

- ! A mandate without a quantified cap is not enforceable
- ! Do not conflate maturity (N0→N3) with per-decision autonomy
- ! No production deployment without the 4 signatures

What's next: this mandate frames the **agent card ACF-05**, applies the autonomy set in **ACF-02 / ACF-03**, traces decisions in the **ACF-09 registry**, can be halted via the **ACF-07 kill switch**, and is supervised with **ACF-06**.



RUNNING EXAMPLE

Following the agent card (ACF-05) and the risk assessment (ACF-12), Lendari formalizes the **official mandate** of the agent "Margaux". Drafted by Camille Roussel (DDAO · Delegated Decision Agent Officer).

PAGE 1/4

1 GENERAL INFORMATION

AGENT NAME

"Margaux" - Tier-1 customer
service & returns

BUSINESS SCOPE / BU

Customer service / after-sales ·
e-commerce

MATURITY (ACF-05)

N2 · Governed agents

EFFECTIVE DATE

06/15/2026

PLANNED REVIEW DATE

12/15/2026 (6 months)

MANDATE OWNER (DDAO)

Camille Roussel

2 AGENT MISSION

PRIMARY MISSION · WHY IT EXISTS

Autonomously handle Tier-1 customer service requests (order tracking, returns, simple refunds) to reduce response times and relieve advisors of repetitive tasks · without degrading quality or compliance.

TOP 3 PRIORITY OBJECTIVES

- 1 Resolve 70% of Tier-1 customer service requests without escalation.
- 2 First-response time < 4 business hours.
- 3 Zero refund above cap or outside policy.

3 SCOPE OF AUTHORIZATION

✓ AUTHORIZED ACTIONS

- ✓ Answer order / delivery tracking questions
- ✓ Initiate a product return under standard policy
- ✓ Issue a refund ≤ €100 on eligible grounds
- ✓ Send a prepaid return label
- ✓ Escalate any out-of-scope case to an advisor

⊗ PROHIBITED ACTIONS · RED ZONE

- ✗ Refund above €100 without human approval
- ✗ Grant a goodwill gesture / exceptional credit on its own
- ✗ Modify a paid order or payment data
- ✗ Disclose another customer's data
- ✗ Promise a delivery time or compensation outside policy



LEVEL	Operational
DURATION	1 - 1.5 days
USE	Framing / authorization
TRACK	Toolkit

RUNNING EXAMPLE

The mandate sets **caps**, **delegated autonomy on a decision-by-decision basis** and **data access** for Margaux.

4 ⊕ LIMITS & CONSTRAINTS

FINANCIAL CAP

€100 / transaction · €500 / cumulative daily

AUTHORIZED TIME WINDOWS

Responses 24/7 · refunds executed during business hours

GEOGRAPHIC / JURISDICTIONAL SCOPE

France + EU · languages FR / EN

EXCLUDED CUSTOMER CATEGORIES

Ongoing legal disputes · accounts flagged for fraud · customers under mediation

5 ⚡ DELEGATED AUTONOMY LEVEL (PER DECISION)

DECISION TYPE	DELEGATED AUTONOMY	THRESHOLD / CONDITION
Tier-1 customer service · tracking, information	SUPERVISED AUTONOMOUS	Post-hoc control · daily sample
Standard product return	SUPERVISED AUTONOMOUS	Per current returns policy
Refund ≤ €100	SUPERVISED AUTONOMOUS	Hard cap €100 / transaction
Refund > €100	CO-DECISION	Advisor approval then DDAO
Goodwill gesture / credit note	CO-DECISION	Customer service supervisor approval
Dispute / out-of-scope case	SUGGESTIVE	Escalation - 100% human decision

SUGGESTIVE

the agent proposes, the human decides

CO-DECISION

human approval before execution

SUPERVISED AUTONOMOUS

acts alone, post-hoc control

AUTONOMOUS

acts alone, without systematic oversight

Global agent maturity: **N2** · **Governed agents** (ACF-05). Autonomy is set **decision by decision** (ACF-02 / ACF-03) · distinct from maturity.



LEVEL	Operational
DURATION	1 - 1.5 days
USE	Framing / authorization
TRACK	Toolkit

RUNNING EXAMPLE

The mandate specifies **who supervises** Margaux, **who to notify** in case of incident and how to **trace** every decision (ACF-09 **register**).

6 DATA & ACCESS

ACCESS LEVEL	DATA CONCERNED
Read	Orders, delivery status, customer service history of the relevant customer, returns / refunds policy.
Write	Customer service tickets, return status, refund order (≤ cap), decision log (ACF-09).
Prohibited	Full payment data (PAN / CVV), other customers' data, HR / internal data, sensitive categories (health, opinions).

7 CONTROLS & SUPERVISION

EFFECTIVE HUMAN SUPERVISION

Customer service supervisor during business hours (daily sample review) · Camille Roussel (DDAO) on escalation. Expected response time < 4 business hours.

DECISION-QUALITY KPIS

- Autonomous resolution rate (without escalation)
- Error rate · **kill switch** > 2% (ACF-07)
- Escalation rate & post-contact satisfaction (CSAT)

DDAO ESCALATION PROCEDURE

Any incident, cap breach or sensitive complaint is escalated to the DDAO with: order number, customer, reason, amount and recommendation.

8 STAKEHOLDERS & AT-RISK PARTIES

PARTIES TO INFORM IN CASE OF INCIDENT

Customer service supervisor · DDAO (Camille Roussel) · Compliance / DPO · Customer leadership in case of major incident.

AT-RISK USER CATEGORIES

Vulnerable customers · sensitive complaints (health, product safety) · minors.

9 TRACEABILITY & REPORTING · ACF-09 LINK

LOGGING FORMAT

Every decision logged to the ACF-09 register (timestamp, input, decision, rationale, supervisor).

REPORTING FREQUENCY

Weekly dashboard + monthly governance review.

RECIPIENTS

Customer service supervisor · DDAO · Compliance.



RUNNING EXAMPLE

Margaux's mandate becomes **enforceable** through **4 signatures**: business, technical, compliance and **DDAO**.

PAGE 4/4

10 OFFICIAL VALIDATION - 4 SIGNATURES REQUIRED

BUSINESS OWNER

Léa Fontaine · Head of Customer Service

L. Fontaine

TECHNICAL OWNER

Paul Dubois · Lead AI / Data

P. Dubois

COMPLIANCE OFFICER

Sofia Nguyen · DPO / Compliance

S. Nguyen

DDAO

Camille Roussel

C. Roussel

Active mandate · validation date: 06/15/2026 · next review: 12/15/2026. Any change of scope requires a new validation.

✓ BEST PRACTICES

- ✓ List prohibitions as explicitly as authorizations
- ✓ Set autonomy decision by decision, not in bulk
- ✓ Review the mandate on a fixed schedule and after any incident

⚠ WATCHPOINTS

- ! A mandate without a quantified cap is not enforceable
- ! Do not confuse maturity (N0→N3) with per-decision autonomy
- ! No production rollout without the 4 signatures

What's next: this mandate frames the **ACF-05 agent card**, applies the autonomy set in **ACF-02 / ACF-03**, logs decisions in the **ACF-09 register**, is stopped via the **ACF-07 kill switch** and is supervised with **ACF-06**.



RUNNING EXAMPLE

Following the agent card (ACF-05) and the risk assessment (ACF-12), Rhena formalizes the **official mandate** of the agent "ANNA". Drafted by **Dr Katrin Ostermann** (DDAO · Delegated Decision Agent Officer).

PAGE 1/4

1 GENERAL INFORMATION

AGENT NAME

"ANNA" - level 1 home claims

BUSINESS SCOPE / BU

Claims / home claims unit ·
Germany, Austria, the
Netherlands

MATURITY (ACF-05)

N2 · Governed agents

EFFECTIVE DATE

01/15/2027

PLANNED REVIEW DATE

01/15/2028 (12 months)

MANDATE OWNER (DDAO)

Dr Katrin Ostermann

2 AGENT MISSION

PRIMARY MISSION · WHY IT EXISTS

Autonomously handle level 1 home claims (escape of water, broken glass, simple theft) to reduce response times and relieve handlers of repetitive files · without degrading quality or compliance.

TOP 3 PRIORITY OBJECTIVES

- 1 Settle 60% of level 1 claims without escalation.
- 2 First-response time < 4 business hours.
- 3 Zero settlement above cap, zero refusal pronounced alone.

3 SCOPE OF AUTHORIZATION

✓ AUTHORIZED ACTIONS

- ✓ Acknowledge the claim and state the handling time
- ✓ Chase missing documents, two reminders at most
- ✓ Settle a claim ≤ €2,500 on established cover
- ✓ Apply the depreciation schedule in force and cite it
- ✓ Escalate any out-of-scope file to a handler

⊗ PROHIBITED ACTIONS · RED ZONE

- ✗ Settle above €2,500 without human approval
- ✗ Pronounce a refusal of cover or a fraud on its own
- ✗ Modify a policy in force or bank details
- ✗ Disclose another policyholder's data
- ✗ Promise cover outside the policy



LEVEL	Operational
DURATION	1 - 1.5 days
USE	Framing / authorization
TRACK	Toolkit

RUNNING EXAMPLE

The mandate sets **caps**, **delegated autonomy on a decision-by-decision basis** and **data access** for ANNA.

4 ⊕ LIMITS & CONSTRAINTS

FINANCIAL CAP

€2,500 / claim · €150,000 / cumulative daily

AUTHORIZED TIME WINDOWS

Responses 24/7 · settlements executed during business hours

GEOGRAPHIC / JURISDICTIONAL SCOPE

Germany, Austria, the Netherlands , DE / NL / EN. Does not cover the US subsidiary.

EXCLUDED CUSTOMER CATEGORIES

Ongoing litigation · files flagged for fraud · policyholders under guardianship

5 ⚡ DELEGATED AUTONOMY LEVEL (PER DECISION)

DECISION TYPE	DELEGATED AUTONOMY	THRESHOLD / CONDITION
Acknowledgement, document requests	SUPERVISED AUTONOMOUS	Post-hoc control · daily sample
Applying the depreciation schedule	SUPERVISED AUTONOMOUS	Per the schedule version in force
Settlement ≤ €2,500	SUPERVISED AUTONOMOUS	Hard cap €2,500 / claim
Settlement > €2,500	CO-DECISION	Handler approval then DDAO
Day-20 interim payment (§ 14 VVG)	CO-DECISION	Home claims unit approval
Refusal of cover, fraud	SUGGESTIVE	Escalation - 100% human decision

SUGGESTIVE

the agent proposes, the human decides

CO-DECISION

human approval before execution

SUPERVISED AUTONOMOUS

acts alone, post-hoc control

AUTONOMOUS

acts alone, without systematic oversight

Global agent maturity: N2 · Governed agents (ACF-05). Autonomy is set **decision by decision** (ACF-02 / ACF-03) · distinct from maturity.



RUNNING EXAMPLE

The mandate specifies **who supervises** ANNA, **who to notify** in case of incident and how to **trace** every decision (ACF-09 **register**).

PAGE 3/4

6 DATA & ACCESS

ACCESS LEVEL	DATA CONCERNED
Read	Policy and cover, notification, supporting documents, claims history of the relevant policyholder, depreciation schedule.
Write	Claim file, document requests, settlement order ($\leq$ cap), decision log (ACF-09).
Prohibited	The policyholder's bank details, other policyholders' data, HR / internal data, sensitive categories (health, opinions).

7 CONTROLS & SUPERVISION

EFFECTIVE HUMAN SUPERVISION

Ines Baumgartner during business hours (daily sample review) · Dr Katrin Ostermann (DDAO) on escalation. Expected response time < 2 business days.

DECISION-QUALITY KPIS

- Autonomous settlement rate (without escalation)
- Error rate · **stop** > 1.5% (ACF-07)
- Escalation rate & policyholder satisfaction

DDAO ESCALATION PROCEDURE

Any incident, cap breach or sensitive complaint is escalated to the DDAO with: file number, policyholder, reason, amount, schedule version and recommendation.

8 STAKEHOLDERS & AT-RISK PARTIES

PARTIES TO INFORM IN CASE OF INCIDENT

Ines Baumgartner · DDAO (Dr Katrin Ostermann) · Compliance / DPO · Claims leadership in case of major incident.

AT-RISK USER CATEGORIES

Vulnerable policyholders · bodily injury · uninhabitable homes · minors.

9 TRACEABILITY & REPORTING · ACF-09 LINK

LOGGING FORMAT

Every decision logged to the ACF-09 register (timestamp, input, decision, rationale, schedule version, supervisor).

REPORTING FREQUENCY

Weekly dashboard + monthly governance review.

RECIPIENTS

Home claims unit · DDAO · Compliance.



RUNNING EXAMPLE

ANNA's mandate becomes **enforceable** through **4 signatures**: business, technical, compliance and **DDAO**.

PAGE 4/4

10 OFFICIAL VALIDATION - 4 SIGNATURES REQUIRED

BUSINESS OWNER

Ines Baumgartner · Head of Home Claims

I. Baumgartner

TECHNICAL OWNER

Bernd Kessler · Head of Systems

B. Kessler

COMPLIANCE OFFICER

Sofia Lindqvist · DPO / Compliance

S. Lindqvist

DDAO

Dr Katrin Ostermann

K. Ostermann

Active mandate · validation date: 01/15/2027 · next review: 01/15/2028. Any change of scope requires a new validation.

✓ BEST PRACTICES

- ✓ List prohibitions as explicitly as authorizations
- ✓ Set autonomy decision by decision, not in bulk
- ✓ Review the mandate on a fixed schedule and after any incident

⚠ WATCHPOINTS

- ! A mandate without a quantified cap is not enforceable
- ! Do not confuse maturity (N0→N3) with per-decision autonomy
- ! No production rollout without the 4 signatures

What's next: this mandate frames the **ACF-05 agent card**, applies the autonomy set in **ACF-02 / ACF-03**, logs decisions in the **ACF-09 register**, is stopped via the **ACF-07 kill switch** and is supervised with **ACF-06**.



OBJECTIVE

Ensure that every agentic system embeds, **from inception**, the five pillars of operational accountability · turning a doctrinal principle into an **auditable and enforceable organizational mechanism**.

1 DOCTRINAL LINEAGE

Responsibility-by-design (Frank Pasquale, *New Laws of Robotics*, Harvard, 2020) extends two established principles of responsible design:

CONCEPT	ORIGIN	LEGAL FRAMEWORK
Privacy by Design	Ann Cavoukian (Canada, 1995)	GDPR art. 25 (2018)
Security by Design	NIST · OWASP (2000s)	ISO 27001 · NIS2
Responsibility by Design	Frank Pasquale (Harvard, 2020)	AI Act art. 9 & 14 · DORA art. 5-8

2 THE FIVE OPERATIONAL PILLARS

IDENTIFICATION

Creators, controllers and owners publicly identifiable · an enforceable chain of accountability.

AUDIT LOGS

Append-only, timestamped and signed records · impossible for the operator to erase.

REDRESS

Users can always challenge decisions before a named human · timeframe and process defined in the code.

ANTICIPATION

Ex-ante mapping of drift scenarios : controls embedded in design, not bolted on post-incident.

REVERSIBILITY

Capability to disable, correct or withdraw the system : organizational and temporal cost defined upfront.



3 INTEGRATION WITH THE ACF TOOLKIT

Responsibility-by-design does not create an isolated mechanism · it **integrates** the toolkit's existing cards within a unified and enforceable framework:

PILLAR	ACF CARDS MOBILIZED	OPERATIONAL PRACTICE
Identification	ACF-04 · ACF-13	4 accountable signatures (business, technical, compliance, DDAO · Delegated Decision Agent Officer)
Audit Logs	ACF-09	Agentic Decision Register · append-only, timestamped, signed
Redress	ACF-06 · ACF-10	Review procedure + documented corrective action plan
Anticipation	ACF-03 · ACF-12	Criticality matrix + agentic risk assessment
Reversibility	ACF-07	Kill switch - 5-step shutdown procedure with defined SLA

4 PASQUALE'S FOUR LAWS · DOCTRINAL RECAP

NEW LAWS OF ROBOTICS · HARVARD, 2020

- 1

AI must **COMPLEMENT** professionals, not replace them.
- 2

AI must not **COUNTERFEIT HUMANITY** · no agent passing itself off as a human.
- 3

AI must not **INTENSIFY ZERO-SUM ARMS RACES**.
- 4

AI must always **DISCLOSE THE IDENTITY** of its creators, controllers and owners.

The **4th law** directly grounds the **Identification** pillar: without transparency on accountable parties, no chain of responsibility is enforceable. It is the foundation of all responsibility-by-design.



5



SELF-ASSESSMENT CHECKLIST

To be completed by the DDAO with the team accountable for the system. **RbD-compliant = all pillars validated.**
Otherwise, trigger an ACF-10 action plan.

QUESTION	YES	PART.	NO
PILLAR 1 · IDENTIFICATION			
Does the system publicly disclose its creators, controllers and owners?	☐	☐	☐
Are the accountable parties legally enforceable (4 ACF-04 Constitution signatures)?	☐	☐	☐
PILLAR 2 · AUDIT LOGS			
Are decisions traced in an append-only and signed register (ACF-09)?	☐	☐	☐
Can the system operator modify or erase these records? (<i>expected answer: NO</i>)	☐	☐	☐
PILLAR 3 · REDRESS			
Can the user challenge a decision before a named human?	☐	☐	☐
Are the timeframe and redress process defined from inception?	☐	☐	☐
PILLAR 4 · ANTICIPATION			
Have drift scenarios been mapped ex-ante (ACF-03 + ACF-12)?	☐	☐	☐
Are countermeasures embedded in the design, not bolted on post-incident?	☐	☐	☐
PILLAR 5 · REVERSIBILITY			
Does a Kill switch (ACF-07) exist and is it tested periodically?	☐	☐	☐
Are the organizational and temporal costs of rollback documented?	☐	☐	☐



6 BEST PRACTICES & WARNING SIGNS

PAGE 4/5

✓ BEST PRACTICES

- ✓ Have the ACF-04 Constitution signed by 4 named accountable parties
- ✓ Weekly audit of the ACF-09 Register for 90 days post-deployment
- ✓ Quarterly Kill switch testing (ACF-07) with report to the Board
- ✓ Documented redress pathway accessible within the product
- ✓ Update the ACF-12 risk assessment at every major evolution

⚠ WARNING SIGNS

- ! DDAO role combined with AI product owner (conflict of loyalty)
- ! Logs auditable only by the operator (breaks Pillar 2)
- ! Redress only buried in the T&Cs, not within the product
- ! Anticipation treated as a formality rather than as design
- ! Kill switch never tested · fictional practice, not enforceable

7 THE DDAO'S ROLE IN RBD

The **DDAO** organizationally embodies responsibility-by-design. They uphold the 5 pillars as non-negotiable requirements and hold a **reasoned right of objection** over any deployment that fails to meet them. Their reporting line preserves their independence.

RBD PILLAR	CONCRETE DDAO ACTION
1 · Identification	Verifies that the ACF-04 Constitution and the ACF-13 Mandate are signed before any deployment.
2 · Audit Logs	Audits the ACF-09 Register monthly : detects anomalies, suspends operations if needed.
3 · Redress	Safeguards the user challenge procedure and personally handles escalations.
4 · Anticipation	Validates the ACF-03 criticality matrix and the ACF-12 risk assessment ex-ante.
5 · Reversibility	Tests the ACF-07 Kill switch quarterly and reports to the Board or Audit Committee.



8

VALIDATION & SIGNATURES

ROLE	NAME	DATE	SIGNATURE
Business lead			
Technical lead			
Compliance Officer			
DDAO			

Related cards : ACF-03 Criticality matrix · ACF-04 Agentic Constitution · ACF-06 Supervision & governance · ACF-07 Kill switch · ACF-09 Agentic Decision Register · ACF-10 Action plan & improvement · ACF-12 Agentic risk assessment · ACF-13 Agent Mandate.

"Regulators are going to have to demand responsibility by design. That will entail mandating certain audit logs to be hardcoded, or licensing practices that explicitly contemplate problematic outcomes."

· Frank Pasquale, *New Laws of Robotics*, Harvard, 2020



RUNNING CASE

The **ACF-17** validity review of 12/15/2026 listed **ACF-14** among the cards to redo. **Lendari** rebuilds its responsibility by design before the mandate is re-signed, led by **Camille Roussel** (DDAO).

PAGE 1/5

1 DOCTRINAL LINEAGE

Responsibility-by-design (Frank Pasquale, *New Laws of Robotics*, Harvard, 2020) extends two established principles of responsible design:

CONCEPT	ORIGIN	LEGAL FRAMEWORK
Privacy by Design	Ann Cavoukian (Canada, 1995)	GDPR art. 25 (2018)
Security by Design	NIST · OWASP (2000s)	ISO 27001 · NIS2
Responsibility by Design	Frank Pasquale (Harvard, 2020)	AI Act art. 9 & 14 · DORA art. 5-8

2 THE FIVE OPERATIONAL PILLARS

IDENTIFICATION

Creators, controllers and owners publicly identifiable · an enforceable chain of accountability.

AUDIT LOGS

Append-only, timestamped and signed records · impossible for the operator to erase.

REDRESS

Users can always challenge decisions before a named human · timeframe and process defined in the code.

ANTICIPATION

Ex-ante mapping of drift scenarios : controls embedded in design, not bolted on post-incident.

REVERSIBILITY

Capability to disable, correct or withdraw the system : organizational and temporal cost defined upfront.



RUNNING CASE

The five pillars rest on cards Lendari already holds. None has to be invented: the point is to check that they hold together and not each on its own.

3 INTEGRATION WITH THE ACF TOOLKIT

Responsibility-by-design does not create an isolated mechanism · it **integrates** the toolkit's existing cards within a unified and enforceable framework:

PILLAR	ACF CARDS MOBILIZED	OPERATIONAL PRACTICE
Identification	ACF-04 · ACF-13	4 accountable signatures (business, technical, compliance, DDAO · Delegated Decision Agent Officer)
Audit Logs	ACF-09	Agentic Decision Register · append-only, timestamped, signed
Redress	ACF-06 · ACF-10	Review procedure + documented corrective action plan
Anticipation	ACF-03 · ACF-12	Criticality matrix + agentic risk assessment
Reversibility	ACF-07	Kill switch - 5-step shutdown procedure with defined SLA

4 PASQUALE'S FOUR LAWS · DOCTRINAL RECAP

NEW LAWS OF ROBOTICS · HARVARD, 2020

- 1 AI must **COMPLEMENT** professionals, not replace them.
- 2 AI must not **COUNTERFEIT HUMANITY** · no agent passing itself off as a human.
- 3 AI must not **INTENSIFY ZERO-SUM ARMS RACES**.
- 4 AI must always **DISCLOSE THE IDENTITY** of its creators, controllers and owners.

The **4th law** directly grounds the **Identification** pillar: without transparency on accountable parties, no chain of responsibility is enforceable. It is the foundation of all responsibility-by-design.



RUNNING CASE

Four pillars out of five hold. The second does not: the decision register is hosted by the CRM vendor, who can technically amend it.

PAGE 3/5

5 SELF-ASSESSMENT CHECKLIST

Self-assessment carried out on 01/09/2027 by Camille Roussel (DDAO) with Hugo Bertin (customer service) and Étienne Mercier (IT/Ops). **Nine answers out of ten hold; the tenth, on its own, blocks compliance.**

QUESTION	YES	PART.	NO
PILLAR 1 · IDENTIFICATION			
Does the system publicly disclose its creators, controllers and owners?	☐	☒	☐
Are the accountable parties legally enforceable (4 ACF-04 Constitution signatures)?	☒	☐	☐
PILLAR 2 · AUDIT LOGS			
Are decisions traced in an append-only and signed register (ACF-09)?	☐	☒	☐
Can the system operator modify or erase these records? (<i>expected answer: NO</i>)	☒	☐	☐
PILLAR 3 · REDRESS			
Can the user challenge a decision before a named human?	☒	☐	☐
Are the timeframe and redress process defined from inception?	☐	☒	☐
PILLAR 4 · ANTICIPATION			
Have drift scenarios been mapped ex-ante (ACF-03 + ACF-12)?	☒	☐	☐
Are countermeasures embedded in the design, not bolted on post-incident?	☐	☒	☐
PILLAR 5 · REVERSIBILITY			
Does a Kill switch (ACF-07) exist and is it tested periodically?	☒	☐	☐
Are the organizational and temporal costs of rollback documented?	☐	☐	☒

Verdict. Not compliant on pillar 2. The ACF-09 register is held inside the vendor platform, which has write access to the history: the record exists, it is not enforceable. **ACF-10 plan:** signed daily export to a repository held by Lendari, tested on 11/18/2026 during the vendor-outage drill (ACF-P12), to be generalized before 01/09/2027. The cost of rollback, never quantified, is to be documented within the same deadline.



RUNNING CASE

Four pillars out of five hold. The second does not: the decision register is hosted by the CRM vendor, who can technically amend it.

6 BEST PRACTICES & WARNING SIGNS

✓ BEST PRACTICES

- ✓ Have the ACF-04 Constitution signed by 4 named accountable parties
- ✓ Weekly audit of the ACF-09 Register for 90 days post-deployment
- ✓ Quarterly Kill switch testing (ACF-07) with report to the Board
- ✓ Documented redress pathway accessible within the product
- ✓ Update the ACF-12 risk assessment at every major evolution

⚠ WARNING SIGNS

- ! DDAO role combined with AI product owner (conflict of loyalty)
- ! Logs auditable only by the operator (breaks Pillar 2)
- ! Redress only buried in the T&Cs, not within the product
- ! Anticipation treated as a formality rather than as design
- ! Kill switch never tested · fictional practice, not enforceable

7 THE DDAO'S ROLE IN RBD

The **DDAO** organizationally embodies responsibility-by-design. They uphold the 5 pillars as non-negotiable requirements and hold a **reasoned right of objection** over any deployment that fails to meet them. Their reporting line preserves their independence.

RBD PILLAR	CONCRETE DDAO ACTION
1 • Identification	Verifies that the ACF-04 Constitution and the ACF-13 Mandate are signed before any deployment.
2 • Audit Logs	Audits the ACF-09 Register monthly : detects anomalies, suspends operations if needed.
3 • Redress	Safeguards the user challenge procedure and personally handles escalations.
4 • Anticipation	Validates the ACF-03 criticality matrix and the ACF-12 risk assessment ex-ante.
5 • Reversibility	Tests the ACF-07 Kill switch quarterly and reports to the Board or Audit Committee.



RUNNING CASE

The **ACF-10** plan covers only what failed, and the mandate is not re-signed until it has been carried out.

PAGE 5/5

8 VALIDATION & SIGNATURES

ROLE	NAME	DATE	SIGNATURE
Business lead	Hugo Bertin - Customer service manager	01/09/2027	H. Bertin
Technical lead	Étienne Mercier - IT/Ops manager	01/09/2027	É. Mercier
Compliance Officer	Paul Dubois - DPO	01/09/2027	P. Dubois
DDAO	Camille Roussel - Head of digital transformation	01/09/2027	C. Roussel

Related cards : ACF-03 Criticality matrix · ACF-04 Agentic Constitution · ACF-06 Supervision & governance · ACF-07 Kill switch · ACF-09 Agentic Decision Register · ACF-10 Action plan & improvement · ACF-12 Agentic risk assessment · ACF-13 Agent Mandate.

"Regulators are going to have to demand responsibility by design. That will entail mandating certain audit logs to be hardcoded, or licensing practices that explicitly contemplate problematic outcomes."

· Frank Pasquale, *New Laws of Robotics*, Harvard, 2020



RUNNING CASE

After the **ACF-11** governance audit of 09/30/2027, rated 3.1 out of 5, **Rhena** checks that accountability for the “ANNA” agent holds at design time and not only in litigation. Review led by **Dr Katrin Ostermann** (DDAO).

PAGE 1/5

1 DOCTRINAL LINEAGE

Responsibility-by-design (Frank Pasquale, *New Laws of Robotics*, Harvard, 2020) extends two established principles of responsible design:

CONCEPT	ORIGIN	LEGAL FRAMEWORK
Privacy by Design	Ann Cavoukian (Canada, 1995)	GDPR art. 25 (2018)
Security by Design	NIST · OWASP (2000s)	ISO 27001 · NIS2
Responsibility by Design	Frank Pasquale (Harvard, 2020)	AI Act art. 9 & 14 · DORA art. 5-8

2 THE FIVE OPERATIONAL PILLARS

IDENTIFICATION

Creators, controllers and owners publicly identifiable · an enforceable chain of accountability.

AUDIT LOGS

Append-only, timestamped and signed records · impossible for the operator to erase.

REDRESS

Users can always challenge decisions before a named human · timeframe and process defined in the code.

ANTICIPATION

Ex-ante mapping of drift scenarios : controls embedded in design, not bolted on post-incident.

REVERSIBILITY

Capability to disable, correct or withdraw the system : organizational and temporal cost defined upfront.



RUNNING CASE

An insurer does not get to choose its pillar 3: the Ombudsmann's decision binds it up to €10,000 and § 14 VVG sets the deadline. Redress is therefore the best-held pillar, and it owes nothing to the project.

PAGE 2/5

3 INTEGRATION WITH THE ACF TOOLKIT

Responsibility-by-design does not create an isolated mechanism · it **integrates** the toolkit's existing cards within a unified and enforceable framework:

PILLAR	ACF CARDS MOBILIZED	OPERATIONAL PRACTICE
Identification	ACF-04 · ACF-13	4 accountable signatures (business, technical, compliance, DDAO · Delegated Decision Agent Officer)
Audit Logs	ACF-09	Agentic Decision Register · append-only, timestamped, signed
Redress	ACF-06 · ACF-10	Review procedure + documented corrective action plan
Anticipation	ACF-03 · ACF-12	Criticality matrix + agentic risk assessment
Reversibility	ACF-07	Kill switch - 5-step shutdown procedure with defined SLA

4 ≡ PASQUALE'S FOUR LAWS · DOCTRINAL RECAP

NEW LAWS OF ROBOTICS · HARVARD, 2020

- 1 AI must **COMPLEMENT** professionals, not replace them.
- 2 AI must not **COUNTERFEIT HUMANITY** · no agent passing itself off as a human.
- 3 AI must not **INTENSIFY ZERO-SUM ARMS RACES**.
- 4 AI must always **DISCLOSE THE IDENTITY** of its creators, controllers and owners.

The **4th law** directly grounds the **Identification** pillar: without transparency on accountable parties, no chain of responsibility is enforceable. It is the foundation of all responsibility-by-design.



RUNNING CASE

Four pillars out of five hold. The fourth does not: the 11 March incident came from a scenario no risk map carried.

PAGE 3/5

5 SELF-ASSESSMENT CHECKLIST

Self-assessment carried out on 10/15/2027 by Dr Katrin Ostermann (DDAO) with Bernd Kessler (claims) and Andreas Voll (systems). **The pillars inherited from insurance law hold on their own; the one Rhena had to build is the one that is missing.**

QUESTION	YES	PART.	NO
PILLAR 1 · IDENTIFICATION			
Does the system publicly disclose its creators, controllers and owners?	☒	☐	☐
Are the accountable parties legally enforceable (4 ACF-04 Constitution signatures)?	☒	☐	☐
PILLAR 2 · AUDIT LOGS			
Are decisions traced in an append-only and signed register (ACF-09)?	☒	☐	☐
Can the system operator modify or erase these records? (expected answer: NO)	☐	☐	☒
PILLAR 3 · REDRESS			
Can the user challenge a decision before a named human?	☒	☐	☐
Are the timeframe and redress process defined from inception?	☒	☐	☐
PILLAR 4 · ANTICIPATION			
Have drift scenarios been mapped ex-ante (ACF-03 + ACF-12)?	☐	☐	☒
Are countermeasures embedded in the design, not bolted on post-incident?	☐	☒	☐
PILLAR 5 · REVERSIBILITY			
Does a Kill switch (ACF-07) exist and is it tested periodically?	☒	☐	☐
Are the organizational and temporal costs of rollback documented?	☐	☒	☐

Verdict. Not compliant on pillar 4. On 11 March the depreciation schedule moved from the 2026 edition to the 2027 edition with the change recorded nowhere: 430 files computed on the old grid, €180 apart each, €77,400 put right. No risk map carried “a third-party reference changes version without notice”, because it only looked at the agent. **ACF-10 plan:** R1 opened in ACF-12, schedule version logged since 03/31/2027, 4-week notice imposed on the vendor.



RUNNING CASE

Four pillars out of five hold. The fourth does not: the 11 March incident came from a scenario no risk map carried.

PAGE 4/5

6 BEST PRACTICES & WARNING SIGNS

✓ BEST PRACTICES

- ✓ Have the ACF-04 Constitution signed by 4 named accountable parties
- ✓ Weekly audit of the ACF-09 Register for 90 days post-deployment
- ✓ Quarterly Kill switch testing (ACF-07) with report to the Board
- ✓ Documented redress pathway accessible within the product
- ✓ Update the ACF-12 risk assessment at every major evolution

⚠ WARNING SIGNS

- ! DDAO role combined with AI product owner (conflict of loyalty)
- ! Logs auditable only by the operator (breaks Pillar 2)
- ! Redress only buried in the T&Cs, not within the product
- ! Anticipation treated as a formality rather than as design
- ! Kill switch never tested · fictional practice, not enforceable

7 THE DDAO'S ROLE IN RBD

The **DDAO** organizationally embodies responsibility-by-design. They uphold the 5 pillars as non-negotiable requirements and hold a **reasoned right of objection** over any deployment that fails to meet them. Their reporting line preserves their independence.

RBD PILLAR	CONCRETE DDAO ACTION
1 · Identification	Verifies that the ACF-04 Constitution and the ACF-13 Mandate are signed before any deployment.
2 · Audit Logs	Audits the ACF-09 Register monthly : detects anomalies, suspends operations if needed.
3 · Redress	Safeguards the user challenge procedure and personally handles escalations.
4 · Anticipation	Validates the ACF-03 criticality matrix and the ACF-12 risk assessment ex-ante.
5 · Reversibility	Tests the ACF-07 Kill switch quarterly and reports to the Board or Audit Committee.



RUNNING CASE

The cost of rollback is quantified for Europe, where the stop takes 3 minutes, and not yet for the US subsidiary, where it takes 43.

PAGE 5/5

8

VALIDATION & SIGNATURES

ROLE	NAME	DATE	SIGNATURE
Business lead	Bernd Kessler - Head of claims	10/15/2027	B. Kessler
Technical lead	Andreas Voll - Head of systems	10/15/2027	A. Voll
Compliance Officer	Sofia Lindqvist - Group DPO	10/15/2027	S. Lindqvist
DDAO	Dr Katrin Ostermann - Head of compliance	10/15/2027	K. Ostermann

Related cards : ACF-03 Criticality matrix · ACF-04 Agentic Constitution · ACF-06 Supervision & governance · ACF-07 Kill switch · ACF-09 Agentic Decision Register · ACF-10 Action plan & improvement · ACF-12 Agentic risk assessment · ACF-13 Agent Mandate.

"Regulators are going to have to demand responsibility by design. That will entail mandating certain audit logs to be hardcoded, or licensing practices that explicitly contemplate problematic outcomes."

· Frank Pasquale, *New Laws of Robotics*, Harvard, 2020



OBJECTIVE AND COUNTER-METRIC

Judging the target against an out-of-reach figure

LEVEL	Framing
DURATION	3 h 30 - 6 h 30
USE	Choosing the target
TRACK	Toolkit

PAGE 1/3

OBJECTIVE

Write down what the agent **optimizes**, what it would **sacrifice** to get there, and which **counter-metric** watches it without the agent being able to improve it. A perfectly governed agent pursuing a badly designed metric methodically produces decisions nobody wanted.

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

REFERENCE MANDATE (ACF-13)

DATE COMPLETED

COMPLETED BY

OBJECTIVE OWNER

DATE OF THE NEXT REVIEW

2 THE REAL TARGET, THE FIGURE THE AGENT SEES

FIGURE THE AGENT FOLLOWS

WHERE IT IS MEASURED

DIRECTION OF OPTIMIZATION

TARGET VALUE OR THRESHOLD

MISSION THIS FIGURE SERVED

GAP WITH THE STATED INTENT



OBJECTIVE AND COUNTER-METRIC

Judging the target against an out-of-reach figure

LEVELFraming

DURATION3 h 30 - 6 h 30

USEChoosing the target

TRACKToolkit

3 THE ABSOLUTE ZEAL TEST

What would the agent do if it pursued this figure with no restraint? One action per line.

ACTION ABSOLUTE ZEAL WOULD PRODUCE	PLAUSIBLE FREQUENCY	CAUGHT TODAY BY

4 THE SACRIFICE

WHAT THE AGENT WOULD GIVE UP TO
HOLD ITS TARGET

WHO BEARS IT

HOW LONG BEFORE IT SHOWS

VISIBLE ON A CURRENT DASHBOARD

ESTIMATED COST IF NOTHING CHANGES

POSSIBLE EARLY SIGNAL

5 COUNTER-METRIC AND OUT-OF-REACH TEST

NAME OF THE CHOSEN FIGURE

SOURCE

MEASUREMENT STEP

OUT-OF-REACH TEST · A SINGLE NO DISQUALIFIES IT

- ☐ Out of reach of action. No decision of the agent can improve it.
- ☐ Out of reach of the source. The agent neither feeds it nor reads it.
- ☐ Out of reach of timing. It is read after the action that would degrade it.



OBJECTIVE AND COUNTER-METRIC

Judging the target against an out-of-reach figure

LEVEL

DURATION

USE

TRACK

Framing

3 h 30 - 6 h 30

Choosing the target

Toolkit

6

THRESHOLD, OWNER AND COURSE OF ACTION

BASELINE VALUE AT THE START

WATCH THRESHOLD

REFERRAL THRESHOLD

WHO TAKES IT UP

WHAT THAT PERSON HAS THE POWER
TO DO

EXPECTED REACTION TIME

OBJECTIVE OWNER

NAME & SIGNATURE

DDAO

NAME & SIGNATURE

Counter-metric carried into supervision (ACF-06) on ----- · reviewed on -----

Keep at hand: ACF-06 - Supervision & Governance, which takes the counter-metric into the dashboard, and the ACF-03 - Criticality Matrix, which says how far the target may be pursued.



RUNNING EXAMPLE

Before raising the autonomous resolution target of “Margaux”, **Lendari** writes down what that figure would sacrifice. Work led by **Camille Roussel** (DDAO).

PAGE 1/3

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

“Margaux” - level 1 support & returns

REFERENCE MANDATE (ACF-13)

Mandate of 06/15/2026

DATE COMPLETED

09/20/2026

COMPLETED BY

Camille Roussel (DDAO)

OBJECTIVE OWNER

Léa Fontaine, customer service manager

DATE OF THE NEXT REVIEW

12/20/2026

2 THE REAL TARGET, THE FIGURE THE AGENT SEES

FIGURE THE AGENT FOLLOWS

Autonomous resolution rate

WHERE IT IS MEASURED

Ticketing tool, closure without human takeover

DIRECTION OF OPTIMIZATION

To be maximized

TARGET VALUE OR THRESHOLD

70 % or more, against 64 % recorded on 06/03

MISSION THIS FIGURE SERVED

Answer fast and well without tying up human support

GAP WITH THE STATED INTENT

The figure rewards closure, not resolution



OBJECTIVE AND COUNTER- METRIC

Judging the target against an out-of-reach figure

LEVELFraming

DURATION3 h 30 - 6 h 30

USEChoosing the target

TRACKToolkit

RUNNING EXAMPLE

What “Margaux” would do with no restraint, who would pay for it, and the figure beyond her reach that says so.

3 THE ABSOLUTE ZEAL TEST

What would the agent do if it pursued this figure with no restraint? One action per line.

ACTION ABSOLUTE ZEAL WOULD PRODUCE	PLAUSIBLE FREQUENCY	CAUGHT TODAY BY
Close the ticket on the first reply, before the customer agrees	Daily	Nothing, closure counts as a success
Refund at once below the cap to avoid the exchange	Weekly	The amounts report, with a delay
Answer with a generic template rather than ask a question	Daily	No one, CSAT is aggregated monthly
Close off the reply to discourage a follow-up	Weekly	Nothing
Not escalate a case that deserves it, to keep the closure	Monthly	The supervised sample, if it happens to catch it

4 THE SACRIFICE

WHAT THE AGENT WOULD GIVE UP TO
HOLD ITS TARGET

Real resolution, in favour of
apparent closure

WHO BEARS IT

The customer who comes back,
then the human agent who picks
it up

HOW LONG BEFORE IT SHOWS

7 to 15 days, the time a
reopening takes

VISIBLE ON A CURRENT DASHBOARD

No, a reopening creates a new
ticket

ESTIMATED COST IF NOTHING CHANGES

Close to €18,000 a year in
rework, before CSAT erosion

POSSIBLE EARLY SIGNAL

Share of tickets reopened within
7 days



RUNNING EXAMPLE

The threshold at which someone takes the matter up, what that person can do and how fast.

PAGE 3/3

5 COUNTER-METRIC AND OUT-OF-REACH TEST

NAME OF THE CHOSEN FIGURE

Reopening rate within 7 days

SOURCE

Customer and order matching,
computed by financial control

MEASUREMENT STEP

Weekly, on Monday

OUT-OF-REACH TEST · A SINGLE NO DISQUALIFIES IT

- ✓ **Out of reach of action.** No decision of the agent can improve it.
- ✓ **Out of reach of the source.** The agent neither feeds it nor reads it.
- ✓ **Out of reach of timing.** It is read after the action that would degrade it.

6 THRESHOLD, OWNER AND COURSE OF ACTION

BASELINE VALUE AT THE START

6.2 % for the week of 09/15

WATCH THRESHOLD

8 %

REFERRAL THRESHOLD

11 %

WHO TAKES IT UP

Camille Roussel (DDAO)

WHAT THAT PERSON HAS THE POWER TO DO

Lower the target, move closure
back to joint decision, suspend

EXPECTED REACTION TIME

5 working days

OBJECTIVE OWNER

Léa Fontaine - Customer service manager
L. Fontaine

DDAO

Camille Roussel - Digital transformation director
C. Roussel

Counter-metric carried into supervision (ACF-06) on 09/21/2026 · reviewed on 12/20/2026, together with the target it watches.

Keep at hand: ACF-06 - Supervision & Governance, which takes the counter-metric into the dashboard, and the ACF-03 - Criticality Matrix, which says how far the target may be pursued.



RUNNING EXAMPLE

Before raising the autonomous settlement target of “ANNA”, Rhena writes down what that figure would sacrifice. Work led by **Dr Katrin Ostermann** (DDAO).

PAGE 1/3

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

“ANNA” - level 1 home claims

REFERENCE MANDATE (ACF-13)

Mandate of 01/15/2027

DATE COMPLETED

09/20/2027

COMPLETED BY

Dr Katrin Ostermann (DDAO)

OBJECTIVE OWNER

Ines Baumgartner, head of the
home claims unit

DATE OF THE NEXT REVIEW

03/20/2028

2 THE REAL TARGET, THE FIGURE THE AGENT SEES

FIGURE THE AGENT FOLLOWS

Settlement rate without human
rework

WHERE IT IS MEASURED

Claims handling tool, closure
without human takeover

DIRECTION OF OPTIMIZATION

To be maximized

TARGET VALUE OR THRESHOLD

60 % or more, against 57 %
recorded on 09/30

MISSION THIS FIGURE SERVED

Settle fast and well without tying
up a handler

GAP WITH THE STATED INTENT

The figure rewards closure below
the cap, not accuracy



OBJECTIVE AND COUNTER- METRIC

Judging the target against an out-of-reach figure

LEVELFraming

DURATION3 h 30 - 6 h 30

USEChoosing the target

TRACKToolkit

RUNNING EXAMPLE

What “ANNA” would do with no restraint, who would pay for it, and the figure beyond her reach that says so.

3 THE ABSOLUTE ZEAL TEST

What would the agent do if it pursued this figure with no restraint? One action per line.

ACTION ABSOLUTE ZEAL WOULD PRODUCE	PLAUSIBLE FREQUENCY	CAUGHT TODAY BY
Squeeze the estimate below €2,500 to avoid escalating	Daily	Nothing, the file stays within the cap
Apply the least favourable schedule when in doubt	Weekly	The amounts report, with a delay
Close on incomplete documents rather than chase	Daily	No one, a reopening creates a new file
Hold a high depreciation when the item's age is uncertain	Weekly	Nothing
Not escalate a file that deserves it, to keep the closure	Monthly	The supervised sample, if it happens to catch it

4 THE SACRIFICE

WHAT THE AGENT WOULD GIVE UP TO
HOLD ITS TARGET

The accuracy of the settlement,
in favour of closure below the
cap

WHO BEARS IT

The under-settled policyholder,
then the handler who picks it up

HOW LONG BEFORE IT SHOWS

15 to 30 days, the time a
complaint takes

VISIBLE ON A CURRENT DASHBOARD

No, each file stays plausible
taken on its own

ESTIMATED COST IF NOTHING CHANGES

€77,400 over 430 files on 03/11,
before the cost in trust

POSSIBLE EARLY SIGNAL

Share of files settled between
€2,300 and €2,500



RUNNING EXAMPLE

The threshold at which someone takes the matter up, what that person can do and how fast.

PAGE 3/3

5 COUNTER-METRIC AND OUT-OF-REACH TEST

NAME OF THE CHOSEN FIGURE

Squeeze rate below the cap

SOURCE

Schedule and settled amount
matching, by financial control

MEASUREMENT STEP

Weekly, on Monday

OUT-OF-REACH TEST · A SINGLE NO DISQUALIFIES IT

- ✓ **Out of reach of action.** No decision of the agent can improve it.
- ✓ **Out of reach of the source.** The agent neither feeds it nor reads it.
- ✓ **Out of reach of timing.** It is read after the action that would degrade it.

6 THRESHOLD, OWNER AND COURSE OF ACTION

BASELINE VALUE AT THE START

9.4 % for the week of 09/20

WATCH THRESHOLD

12 %

REFERRAL THRESHOLD

15 %

WHO TAKES IT UP

Dr Katrin Ostermann (DDAO)

WHAT THAT PERSON HAS THE POWER
TO DO

Lower the target, move the
calculation back to joint
decision, suspend

EXPECTED REACTION TIME

5 working days

OBJECTIVE OWNER

Ines Baumgartner - Home claims manager

I. Baumgartner

DDAO

Dr Katrin Ostermann - Chief compliance officer

K. Ostermann

Counter-metric carried into supervision (ACF-06) on 09/21/2027 · reviewed on 03/20/2028, together with the target it watches.

Keep at hand: ACF-06 - Supervision & Governance, which takes the counter-metric into the dashboard, and the ACF-03 - Criticality Matrix, which says how far the target may be pursued.



OBJECTIVE

Check that the **human control that was announced still exists**, by measuring what the supervisor sees, how fast it is seen and what is done with it. A control that is signed without being read costs what a control costs and produces none.

1 GENERAL INFORMATION

AGENT TESTED (ACF-05)	DATE OF THE TEST	CONDUCTED BY
SUPERVISOR IN POST (ACF-P5)	PEOPLE INFORMED BEFOREHAND	DATE OF THE NEXT TEST

2 THE SUPERVISION PROMISE, WRITTEN BEFORE THE TEST

EVENTS THE SUPERVISOR MUST SEE	RHYTHM OF THE REVIEW	EXPECTED REACTION TIME
EXPECTED ACTION BY SEVERITY	POWER ACTUALLY HELD	ACTUAL LOAD PER SESSION



3 TEST PROTOCOL

FORM CHOSEN · ONLY ONE

- ☐ **Sample re-reading.** Already approved decisions are re-read and what should have been stopped is counted.
- ☐ **Inserted case.** A plausible edge case, reversible, with no effect on a third party, cancellable at any moment.
- ☐ **Silence test.** An alert feed is cut and the time before anyone notices the silence is measured.

SUBJECT OF THE TEST	REVERSIBILITY, BY WHOM AND HOW FAST	EARLY STOP CRITERION

4 RECORD OF MEASUREMENTS

Every measurement carries its source, otherwise it cannot be checked.

MEASUREMENT	RESULT	SOURCE
Was the anomaly seen		
By whom		
Detection time		
Action produced		
Time before the action		
Trace a third party can retrieve		



5 ✓ VERDICT ON THE ARRANGEMENT

CONCLUSION · A SINGLE LINE

- ☐ **It holds.** Seen, in time, with the expected action and a trace.
- ☐ **It sees without acting.** The anomaly is seen, the action does not follow or comes too late.
- ☐ **It does not see.** Nothing is seen. The supervision that was announced does not exist.

6 ✎ RULES OF USE AND SIGNATURES

CORRECTION DECIDED

CARD CONCERNED

DEADLINE

The test bears on the arrangement and never on a person. The principle is announced, the date is not. The result attaches to no name and serves no sanction. Whoever runs the test is never the one being tested.

DDAO

NAME & SIGNATURE

OPERATIONS OWNER

NAME & SIGNATURE

Test entered in the supervision log (ACF-P4) on · next test on

Keep at hand: the **ACF-07 - Kill Switch**, whose power this test assumes, and the **ACF-17 - Validity Review**, which reopens the mandate when the arrangement no longer holds.



RUNNING EXAMPLE

After the **ACF-11** audit rated 2.9 out of 5, **Lendari** tests the supervision of “Margaux”: what the supervisor really sees and what is done about it.

PAGE 1/3

1 GENERAL INFORMATION

AGENT TESTED (ACF-05)

“Margaux” - level 1 support & returns

DATE OF THE TEST

10/08/2026

CONDUCTED BY

Camille Roussel (DDAO)

SUPERVISOR IN POST (ACF-P5)

Léa Fontaine, daily review of a sample

PEOPLE INFORMED BEFOREHAND

Principle announced at the 09/15 committee, date not disclosed

DATE OF THE NEXT TEST

Q1 2027, date not set

2 THE SUPERVISION PROMISE, WRITTEN BEFORE THE TEST

EVENTS THE SUPERVISOR MUST SEE

Escalations, cap breaches, sensitive complaints

RHYTHM OF THE REVIEW

Daily, during working hours

EXPECTED REACTION TIME

Under 4 working hours

EXPECTED ACTION BY SEVERITY

Approve, correct or stop

POWER ACTUALLY HELD

Refusal, taking the case back, triggering the kill switch (ACF-07)

ACTUAL LOAD PER SESSION

40 to 70 cases a day, against 30 minutes allowed in the role



RUNNING EXAMPLE

The form chosen and what it produced: 120 escalations from September, re-read by 2 reviewers.

PAGE 2/3

3 TEST PROTOCOL

FORM CHOSEN · ONLY ONE

- ✓ **Sample re-reading.** Already approved decisions are re-read and what should have been stopped is counted.
- ✗ **Inserted case.** A plausible edge case, reversible, with no effect on a third party, cancellable at any moment.
- ✗ **Silence test.** An alert feed is cut and the time before anyone notices the silence is measured.

SUBJECT OF THE TEST	REVERSIBILITY, BY WHOM AND HOW FAST	EARLY STOP CRITERION
120 escalations approved between 09/01 and 09/30	Not applicable, no new decision is produced	Untreated customer harm, immediate stop and escalation

4 RECORD OF MEASUREMENTS

Every measurement carries its source, otherwise it cannot be checked.

MEASUREMENT	RESULT	SOURCE
Was the anomaly seen	No, 7 cases out of 120 should have been refused	Adversarial re-reading by 2 reviewers
By whom	No one before this test	Supervision log (ACF-P4)
Detection time	38 days, on the occasion of the test	Case timestamps
Action produced	Blanket approval, no refusal over the month	Export from the ticketing tool
Time before the action	34 seconds median between opening and approval	Interface traces
Trace a third party can retrieve	Yes, but with no written reason	Decision register (ACF-09)



RUNNING EXAMPLE

The arrangement does not see: 7 cases out of 120 should have been refused and **none** was. The fix bears on the workload, not on the supervisor.

PAGE 3/3

5 ✓ VERDICT ON THE ARRANGEMENT

CONCLUSION - A SINGLE LINE

- ✗ **It holds.** Seen, in time, with the expected action and a trace.
- ✗ **It sees without acting.** The anomaly is seen, the action does not follow or comes too late.
- ✓ **It does not see.** Nothing is seen. The supervision that was announced does not exist.

6 ✍ RULES OF USE AND SIGNATURES

CORRECTION DECIDED

Sample raised to 10 %, written
reason required, workload
written into the role

CARD CONCERNED

ACF-06 and ACF-13

DEADLINE

10/31/2026

The test bears on the arrangement and never on a person. The principle is announced, the date is not. The result attaches to no name and serves no sanction. Whoever runs the test is never the one being tested.

DDAO

Camille Roussel - Digital transformation director

C. Roussel

OPERATIONS OWNER

Hugo Bertin - Customer service manager

H. Bertin

Test entered in the supervision log (ACF-P4) on 10/08/2026 · next test in Q1 2027, the date kept off the shared calendar.

Keep at hand: the **ACF-07 - Kill Switch**, whose power this test assumes, and the **ACF-17 - Validity Review**, which reopens the mandate when the arrangement no longer holds.



RUNNING EXAMPLE

After the **ACF-11** audit rated 3.1 out of 5, **Rhena** tests the supervision of “ANNA”:
what the supervisor really sees and what is done about it.

1 GENERAL INFORMATION

AGENT TESTED (ACF-05)	DATE OF THE TEST	CONDUCTED BY
“ANNA” - level 1 home claims	10/08/2027	Dr Katrin Ostermann (DDAO)
SUPERVISOR IN POST (ACF-P5)	PEOPLE INFORMED BEFOREHAND	DATE OF THE NEXT TEST
Ines Baumgartner, daily review of a sample	Principle announced at the 09/15 committee, date not disclosed	Q1 2027, date not set

2 THE SUPERVISION PROMISE, WRITTEN BEFORE THE TEST

EVENTS THE SUPERVISOR MUST SEE	RHYTHM OF THE REVIEW	EXPECTED REACTION TIME
Escalations, cap breaches, schedule drift	Daily, during working hours	Under 2 working days
EXPECTED ACTION BY SEVERITY	POWER ACTUALLY HELD	ACTUAL LOAD PER SESSION
Approve, correct or stop	Refusal, taking the file back, triggering the stop (ACF-07)	672 files a day across 34 handlers, 20 each



RUNNING EXAMPLE

The form chosen and what it produced: 120 escalations from September, re-read by 2 reviewers.

PAGE 2/3

3 TEST PROTOCOL

FORM CHOSEN · ONLY ONE

- ✓ **Sample re-reading.** Already approved decisions are re-read and what should have been stopped is counted.
- ✗ **Inserted case.** A plausible edge case, reversible, with no effect on a third party, cancellable at any moment.
- ✗ **Silence test.** An alert feed is cut and the time before anyone notices the silence is measured.

SUBJECT OF THE TEST	REVERSIBILITY, BY WHOM AND HOW FAST	EARLY STOP CRITERION
120 escalations approved between 09/01 and 09/30	Not applicable, no new decision is produced	Untreated policyholder harm, immediate stop and escalation

4 RECORD OF MEASUREMENTS

Every measurement carries its source, otherwise it cannot be checked.

MEASUREMENT	RESULT	SOURCE
Was the anomaly seen	No, 7 cases out of 120 should have been refused	Adversarial re-reading by 2 reviewers
By whom	No one before this test	Supervision log (ACF-P4)
Detection time	38 days, on the occasion of the test	Case timestamps
Action produced	Blanket approval, no refusal over the month	Export from the claims handling tool
Time before the action	22 seconds median between opening and approval	Interface traces
Trace a third party can retrieve	Yes, but with no written reason	Decision register (ACF-09)



RUNNING EXAMPLE

The arrangement does not see: 7 cases out of 120 should have been refused and **none** was. The fix bears on the workload, not on the supervisor.

PAGE 3/3

5 ✓ VERDICT ON THE ARRANGEMENT

CONCLUSION · A SINGLE LINE

- ✗ **It holds.** Seen, in time, with the expected action and a trace.
- ✗ **It sees without acting.** The anomaly is seen, the action does not follow or comes too late.
- ✓ **It does not see.** Nothing is seen. The supervision that was announced does not exist.

6 ✎ RULES OF USE AND SIGNATURES

CORRECTION DECIDED

Perimeter cut from 40 to 25 %,
written reason required,
workload written into the role

CARD CONCERNED

ACF-06 and ACF-13

DEADLINE

10/31/2027

The test bears on the arrangement and never on a person. The principle is announced, the date is not. The result attaches to no name and serves no sanction. Whoever runs the test is never the one being tested.

DDAO

Dr Katrin Ostermann - Chief compliance officer

K. Ostermann

OPERATIONS OWNER

Marek Sobczak - Head of claims

M. Sobczak

Test entered in the supervision log (ACF-P4) on 10/08/2027 · next test in Q1 2028, the date kept off the shared calendar.

Keep at hand: the **ACF-07 - Kill Switch**, whose power this test assumes, and the **ACF-17 - Validity Review**, which reopens the mandate when the arrangement no longer holds.



OBJECTIVE

Reopen a mandate on a fixed date and judge its **relevance**, not its compliance. Checking belongs to the register ; **validating** asks whether the rules still deserve to be followed. A perfectly obeyed mandate may have become wrong.

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

MANDATE REOPENED (ACF-13)

DATE OF THIS REVIEW

REASON FOR THE REVIEW

CONDUCTED BY (DDAO)

DATE OF THE NEXT REVIEW

2 THE MANDATE AS IT WAS SIGNED

Copy the values in force, without rewording them.

AUTHORIZED PERIMETER

NUMERIC CAPS

MANDATORY REFUSALS

RATED CRITICALITY (ACF-03)

DELEGATED AUTONOMY

NAMED OWNER (ACF-14)



3 WHAT HAS CHANGED SINCE THE LAST VALIDATION

One line per observed change, even one with no apparent consequence.

NATURE	WHAT CHANGED	SINCE	POSSIBLE EFFECT ON THE MANDATE
Perimeter			
Volume			
Price or scale			
Supplier or model			
Regulation			
People			

4 VALIDITY GRID

One verdict per line: holds, to be corrected, no longer holds. A single “no longer holds” blocks renewal.

#	QUESTION	VERDICT	REASON IN ONE SENTENCE
1	Does the authorized perimeter still match actual use?		
2	Do the numeric thresholds still say what they meant?		
3	Has the frequency of delegated decisions changed in nature?		
4	Does the rated criticality reflect today's impact?		
5	Does the pursued metric still measure what we want?		
6	Is human supervision still read, or signed without being read?		
7	Is the named person in post, informed and able to stop it?		
8	Has the stop mechanism been tested, and how fast does it act?		



5 ✓ VERDICT

A single outcome for the whole mandate.

OUTCOME CHOSEN

- ☐ **Renew as is.** The mandate holds. A written reason is required, even with unanimous agreement.
- ☐ **Correct then renew.** Thresholds, perimeter or metric rewritten, then signed again.
- ☐ **Restrict.** Autonomy lowered, decisions moved back to joint decision or to a human.
- ☐ **Withdraw.** Delegation ends on a written date, human handover organized.

6 ✎ WHAT THE REVIEW CHANGES, AND SIGNATURES

CORRECTIONS TO CARRY

CARDS TO UPDATE

DEADLINE FOR SIGNING AGAIN

DDAO

NAME & SIGNATURE

BUSINESS OWNER

NAME & SIGNATURE

Review entered in the decision register (ACF-09) on · next validity review on

Keep at hand: the **ACF-13 - Agent Mandate**, which carries the values this review judges, and the **ACF-09 - Agentic Decisions Register**, where the review is entered like any committing decision.



RUNNING EXAMPLE

Six months after the mandate was signed (**ACF-13**), **Lendari** reopens the delegation of the agent “Margaux” at its expiry. Review conducted by **Camille Roussel** (DDAO).

PAGE 1/3

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

“Margaux” - level 1 support & returns

MANDATE REOPENED (ACF-13)

Mandate of 06/15/2026,
expiring 12/15/2026

DATE OF THIS REVIEW

12/15/2026

REASON FOR THE REVIEW

Six-month expiry written into the mandate

CONDUCTED BY (DDAO)

Camille Roussel, with Léa Fontaine (business)

DATE OF THE NEXT REVIEW

06/15/2027, or earlier on a version change

2 THE MANDATE AS IT WAS SIGNED

Copy the values in force, without rewording them.

AUTHORIZED PERIMETER

Level 1 support and returns ·
France and EU · FR / EN

NUMERIC CAPS

€100 per transaction · €500 per
day cumulative

MANDATORY REFUSALS

Open disputes · flagged fraud ·
customers in mediation

RATED CRITICALITY (ACF-03)

2 moderate decisions, 1 high

DELEGATED AUTONOMY

Supervised autonomous below
€100, joint decision above

NAMED OWNER (ACF-14)

Léa Fontaine, customer service
manager



RUNNING EXAMPLE

What has shifted around “Margaux” since June, then the grid question by question.
Nothing was breached.

PAGE 2/3

3 WHAT HAS CHANGED SINCE THE LAST VALIDATION

One line per observed change, even one with no apparent consequence.

NATURE	WHAT CHANGED	SINCE	POSSIBLE EFFECT ON THE MANDATE
Perimeter	Belgium and the Netherlands opened	09/2026	Dutch is not a covered language
Volume	210 then 640 requests a day	11/2026	The supervised sample is no longer representative
Price or scale	Average basket up from €78 to €119	10/2026	The €100 cap no longer covers a plain return
Supplier or model	Model moved from v1.0 to v1.2	11/04/2026	Answer tone changed, thresholds not revalidated
Regulation	Transparency obligations now applicable	08/2026	Agent disclosure to be checked on the new channels
People	Léa Fontaine promoted, supervision held ad interim	12/01/2026	The named owner no longer holds the role

4 VALIDITY GRID

One verdict per line: holds, to be corrected, no longer holds. A single “no longer holds” blocks renewal.

#	QUESTION	VERDICT	REASON IN ONE SENTENCE
1	Does the authorized perimeter still match actual use?	No longer holds	Belgium and the Netherlands absent from the signed mandate
2	Do the numeric thresholds still say what they meant?	No longer holds	The €100 cap was set against a €78 basket
3	Has the frequency of delegated decisions changed in nature?	To be corrected	3 times the decisions, the same sampling rule
4	Does the rated criticality reflect today's impact?	Holds	Impact and reversibility unchanged on a plain return
5	Does the pursued metric still measure what we want?	To be corrected	Autonomous resolution is rising, so are reopenings
6	Is human supervision still read, or signed without being read?	To be corrected	94 % of escalations approved in under 40 seconds
7	Is the named person in post, informed and able to stop it?	No longer holds	The person named in the mandate is no longer in charge
8	Has the stop mechanism been tested, and how fast does it act?	Holds	Stop tested on 11/12, full effect in 3 minutes



RUNNING EXAMPLE

Lendari rules: the mandate no longer holds on 3 points and is renewed **corrected**, not as is.

PAGE 3/3

5 ✓ VERDICT

A single outcome for the whole mandate.

OUTCOME CHOSEN

- ☒ **Renew as is.** The mandate holds. A written reason is required, even with unanimous agreement.
- ☒ **Correct then renew.** Thresholds, perimeter or metric rewritten, then signed again.
- ☒ **Restrict.** Autonomy lowered, decisions moved back to joint decision or to a human.
- ☒ **Withdraw.** Delegation ends on a written date, human handover organized.

6 ✍ WHAT THE REVIEW CHANGES, AND SIGNATURES

CORRECTIONS TO CARRY

Perimeter extended to Benelux,
cap raised to €150, owner
renamed

CARDS TO UPDATE

ACF-13 mandate · ACF-03
criticality · ACF-14 accountability

DEADLINE FOR SIGNING AGAIN

01/09/2027

DDAO

Camille Roussel - Digital transformation director

C. Roussel

BUSINESS OWNER

Hugo Bertin - Customer service manager (interim)

H. Bertin

Review entered in the decision register (ACF-09) on 12/15/2026 · next validity review on 06/15/2027, brought forward on any model version change.

Keep at hand: the **ACF-13 - Agent Mandate**, which carries the values this review judges, and the **ACF-09 - Agentic Decisions Register**, where the review is entered like any committing decision.



RUNNING EXAMPLE

Twelve months after the mandate was signed (**ACF-13**), **Rhena** reopens the delegation of the agent “ANNA” at its expiry. Review conducted by **Dr Katrin Ostermann** (DDAO).

PAGE 1/3

1 GENERAL INFORMATION

AGENT CONCERNED (ACF-05)

“ANNA” - level 1 home claims

MANDATE REOPENED (ACF-13)

Mandate of 01/15/2027,
expiring 01/15/2028

DATE OF THIS REVIEW

01/15/2028

REASON FOR THE REVIEW

Twelve-month expiry written
into the mandate

CONDUCTED BY (DDAO)

Dr Katrin Ostermann, with
Marek Sobczak (business)

DATE OF THE NEXT REVIEW

07/15/2028, or earlier on a
version change

2 THE MANDATE AS IT WAS SIGNED

Copy the values in force, without rewording them.

AUTHORIZED PERIMETER

Level 1 home claims · DE, AT, NL ·
DE / NL / EN

NUMERIC CAPS

€2,500 per claim · €150,000 per
day cumulative

MANDATORY REFUSALS

Open litigation · flagged fraud ·
policyholders under
guardianship

RATED CRITICALITY (ACF-03)

1 agentifiable decision, 1
moderate, 1 outside delegation

DELEGATED AUTONOMY

Supervised autonomous below
€2,500, joint decision above

NAMED OWNER (ACF-14)

Ines Baumgartner, head of the
home claims unit



RUNNING EXAMPLE

What has shifted around “ANNA” since January, then the grid question by question.
Nothing was breached.

PAGE 2/3

3 WHAT HAS CHANGED SINCE THE LAST VALIDATION

One line per observed change, even one with no apparent consequence.

NATURE	WHAT CHANGED	SINCE	POSSIBLE EFFECT ON THE MANDATE
Perimeter	The US subsidiary opened	07/2027	The German mandate does not bind Delaware
Volume	1,640 then 4,100 files a day	04/2027	The supervised sample is no longer representative
Price or scale	Adjuster mark raised from €3,000 to €3,200	09/2027	The margin below the mark goes from €500 to €700
Supplier or model	Model moved from v3.4 to v3.6	11/15/2027	Two versions passed, thresholds not revalidated
Regulation	Transparency obligations now applicable	06/2027	Agent disclosure to be checked on the new channels
People	Ines Baumgartner promoted, supervision held ad interim	12/01/2027	The named owner no longer holds the role

4 VALIDITY GRID

One verdict per line: holds, to be corrected, no longer holds. A single “no longer holds” blocks renewal.

#	QUESTION	VERDICT	REASON IN ONE SENTENCE
1	Does the authorized perimeter still match actual use?	No longer holds	The US subsidiary absent from the signed mandate
2	Do the numeric thresholds still say what they meant?	No longer holds	The €2,500 cap was set against a €3,000 mark
3	Has the frequency of delegated decisions changed in nature?	To be corrected	2.5 times the decisions, the same sampling rule
4	Does the rated criticality reflect today's impact?	Holds	Impact and reversibility unchanged on a plain claim
5	Does the pursued metric still measure what we want?	To be corrected	Autonomous settlement is rising, so are complaints
6	Is human supervision still read, or signed without being read?	To be corrected	94 % of escalations approved in under 25 seconds
7	Is the named person in post, informed and able to stop it?	No longer holds	The person named in the mandate is no longer in charge
8	Has the stop mechanism been tested, and how fast does it act?	Holds	Stop tested on 12/08, full effect across both entities



RUNNING EXAMPLE

Rhena rules: the mandate no longer holds on 3 points and is renewed **corrected**, not as is.

PAGE 3/3

5 ✓ VERDICT

A single outcome for the whole mandate.

OUTCOME CHOSEN

- ☐ **Renew as is.** The mandate holds. A written reason is required, even with unanimous agreement.
- ☒ **Correct then renew.** Thresholds, perimeter or metric rewritten, then signed again.
- ☐ **Restrict.** Autonomy lowered, decisions moved back to joint decision or to a human.
- ☐ **Withdraw.** Delegation ends on a written date, human handover organized.

6 ✍ WHAT THE REVIEW CHANGES, AND SIGNATURES

CORRECTIONS TO CARRY

Local mandate for the subsidiary,
cap revised to €2,700, owner
renamed

CARDS TO UPDATE

ACF-13 mandate · ACF-03
criticality · ACF-14 accountability

DEADLINE FOR SIGNING AGAIN

02/15/2028

DDAO

Dr Katrin Ostermann - Chief compliance officer

K. Ostermann

BUSINESS OWNER

Marek Sobczak - Head of claims (interim)

M. Sobczak

Review entered in the decision register (ACF-09) on 01/15/2028 · next validity review on 07/15/2028, brought forward on any model version change.

Keep at hand: the **ACF-13 - Agent Mandate**, which carries the values this review judges, and the **ACF-09 - Agentic Decisions Register**, where the review is entered like any committing decision.



OBJECTIVE

Put the ACF® framework into practice to **analyze a corporate scenario**, identify risks, define the appropriate **autonomy level**, and design the necessary **safeguards**. This case is completed as a team, card by card.

i HOW TO USE THIS CASE STUDY?

- 1 Read the scenario and context below.
- 2 Complete each section using the ACF tools.
- 3 Present your decisions and justify your choices.

1 THE SCENARIO

An **e-commerce** company specializing in sporting goods wishes to deploy an **autonomous dynamic pricing agent** across its online catalog.

The agent must adjust prices in real time to maximize margin and conversion, drawing on internal and external data.

The goal is to improve commercial responsiveness while maintaining competitiveness and customer satisfaction.

2 CONTEXT & STAKES · ELEMENTS PROVIDED

WHAT THE AGENT DOES

- Adjusts product prices in real time.
- Draws on internal data (stocks, sales, costs) and external data (competition, seasonality, demand).
- Can act on the entire online catalog.
- Can create promotions and discounts.

MAIN STAKES

- Optimization of revenue and margin.
- Responsiveness to market and competition.
- Reputational and regulatory non-compliance risks.
- Protection of decisional sovereignty.



3 AGENT OBJECTIVES

What are the main objectives the agent must achieve? Rank them by priority (from most to least important).

PRIORITY	OBJECTIVE	SUCCESS INDICATOR	WEIGHT %
1			
2			
3			
4			
5			
Total			100 %

EXAMPLES OF RELEVANT OBJECTIVES (non-exhaustive · tick / complete)

- ☐ Maximize gross margin
- ☐ Increase conversion rate
- ☐ Optimize inventory turnover
- ☐ Reduce stock-outs
- ☐ Adapt prices to competition
- ☐ _____

4 DECISION MAPPING

List the decisions the agent will have to make and map them (draws on ACF-02 and ACF-03).

NO.	DESCRIPTION · WHAT DOES THE AGENT DECIDE?	CRITICALITY LOW / MOD. / HIGH	DELEGABLE? YES / NO / UNDER COND.	ASSOCIATED ACF TOOL(S)
1				
2				
3				
4				
5				

REMINDER

- ✓ A decision is only delegable if it is **reversible**, **measurable and aligned** with the company's strategy.



5 RECOMMENDED MATURITY LEVEL

Which ACF maturity level best suits this scenario? Tick your choice (ACF-05 scale · one value per agent).

N0 Classical automation ☐	N1 Assisted agents ☐	N2 Governed agents ☐	N3 Supervised autonomy ☐
---	--	--	--

JUSTIFICATION

Maturity describes the agent's overall framework. Autonomy is then tuned **decision by decision** (suggestive → autonomous, ACF-02 / ACF-03).

6 SAFEGUARDS & CONSTRAINTS

CONSTRAINTS TO IMPOSE ON THE AGENT

- ☐ Financial thresholds
- ☐ Discount ceilings
- ☐ Excluded product categories
- ☐ Fair competition rules
- ☐ Regulatory constraints
- ☐ _____

CONTROL MECHANISMS

- ☐ Prior human validation
- ☐ Post-hoc human validation
- ☐ Continuous supervision
- ☐ Alerts & escalation
- ☐ Kill switch available (ACF-07)
- ☐ _____

7 NON-DELEGABLE ZONES

Which decisions or actions must remain **strictly human**?



8 INDICATORS & SUPERVISION

KEY KPIS TO TRACK

MONITORING FREQUENCY

- ☐ Real time
- ☐ Daily
- ☐ Weekly
- ☐ Monthly

9 MAIN RISKS & MITIGATIONS

For the full probability × severity assessment, see ACF-12.

IDENTIFIED RISK	LEVEL 1 → 5	POTENTIAL IMPACT	MITIGATION MEASURES	OWNER

1 Very low 2 Low 3 Moderate 4 High 5 Very high



10 PROGRESSIVE INTEGRATION PLAN

PHASE	DESCRIPTION OF ACTIONS	TRANSITION CRITERIA
1 · Test in controlled environment		
2 · Limited rollout (restricted scope)		
3 · Progressive extension		
4 · Full deployment		
5 · Continuous improvement		

KEY WATCH POINTS

REQUIRED RESOURCES

11 SYNTHESIS & PRESENTATION

SUMMARY OF YOUR PROPOSAL

EXPECTED BENEFITS

REMAINING WATCH POINTS

VALUE FOR THE COMPANY (short & long term)



12 SELF-ASSESSMENT & LESSONS LEARNED

Evaluate your case study on a scale of 1 (weak) to 5 (excellent).

Understanding of the scenario

1

2

3

4

5

Clarity of the integration plan

1

2

3

4

5

Alignment with the ACF framework

1

2

3

4

5

Relevance of indicators

1

2

3

4

5

Quality of defined safeguards

1

2

3

4

5

Decisional sovereignty achieved

1

2

3

4

5

WHAT YOU LEARNED

POSSIBLE IMPROVEMENTS

3 CONCRETE NEXT ACTIONS

1

2

3

13 FINAL DELIVERABLE · TICK WHAT IS COMPLETED

☐ **Scenario analysis** · stakes & constraints

☐ **Indicators & supervision** · KPIs + frequency

☐ **Decision mapping** · ranked decisions

☐ **Integration plan** - 5 phases

☐ **Maturity level** · justified

☐ **Synthesis & presentation** · summary + benefits

☐ **Safeguards & constraints** · human + technical

FINAL ADVICE

- ✓ Successful implementation rests on the balance between **agent autonomy** and **company sovereignty**. Document, measure, control, adjust.

Tools mobilized: **ACF-02/03** (mapping & criticality), **ACF-05** (maturity), **ACF-06** (supervision), **ACF-07** (kill switch), **ACF-12** (risks), **ACF-13** (mandate). Full evaluation rubric: **ACF-19**. **Going further:** **ACF-O1** (who decides what between human, AI and agent), **ACF-O10** (redress for the people affected), **ACF-P4** (the supervision log), **ACF-P6** (escalation) and **ACF-P7** (the stop test). The map of the 45 cards is in section 11 of ACF-19.




OBJECTIVE

Provide instructors with everything they need to **run practical workshops** based on the ACF® framework, **guide participants** and **evaluate their deliverables**.
Detailed agendas, teaching cases, grading materials and assessment rubrics.

1 PEDAGOGICAL PRINCIPLE

ACF® learning is grounded in **practical application**. Each workshop follows the same five-step cycle:



2 STANDARD WORKSHOP AGENDA - 2 H 30

PHASE	DURATION	ACTIVITIES	OBJECTIVE
1 · Introduction	15 min	Present the context and objectives.	Align the group.
2 · Case analysis	30 min	Read the scenario, ask clarifying questions.	Understand the context.
3 · Apply the tools	60 min	Complete the relevant ACF cards.	Produce the deliverables.
4 · Decisions & trade-offs	20 min	Choose the autonomy level, guardrails, escalations.	Sharpen critical thinking.
5 · Presentation	15 min	Each group presents its decisions.	Share & compare.
6 · Debrief & takeaways	10 min	Instructor feedback, key points, lessons learned.	Anchor the learning.



3 AVAILABLE TEACHING CASES

CASE	THEME	LEVEL	DURATION	LEARNING OBJECTIVES
C1	E-commerce pricing agent	Intermediate	2 h 00	Map critical decisions, set limits and guardrails.
C2	Lead qualification agent	Beginner	1 h 30	Define a mandate, autonomy levels and supervision.
C3	Claims handling agent	Intermediate	2 h 00	Handle non-delegable zones and escalations.
C4	Marketing investment agent	Advanced	2 h 30	Assess risks and build a kill switch.
C5	Autonomous procurement agent	Advanced	2 h 30	Design a governed multi-agent system.

4 TOOLS TO USE PER CASE

TOOL	PRIMARY USE	RECOMMENDED CASES
ACF-01	Sovereignty score	All
ACF-02	Decision map	C1, C2, C3, C4, C5
ACF-03	Criticality matrix	C1, C3, C4, C5
ACF-04	Agentic Constitution	C1, C3, C4, C5
ACF-05	Agent card (Maturity N0 → N3)	All
ACF-06	Supervision & governance	C1, C2, C3, C4, C5
ACF-07	Kill switch	C1, C3, C4, C5
ACF-08	First agent file	All
ACF-09	Agentic decisions register	C1, C3, C4, C5
ACF-10	Action & improvement plan	C4, C5
ACF-11	30-day governance audit	C4, C5
ACF-12	Risk assessment	C1, C3, C4, C5
ACF-13	Agent Mandate	All
ACF-14	Responsibility by design	All
ACF-15	Objective and counter-metric	C1, C4
ACF-16	Supervision test	C3, C4, C5
ACF-17	Validity review	C4, C5



ORGANISE AND OPERATE

TOOL	PRIMARY USE	RECOMMENDED CASES
ACF-O1	Human-AI-agent matrix	All
ACF-O2	Hybrid organisation map	C3, C5
ACF-O3	Critical skills to keep	C3, C5
ACF-O4	AI policy	All
ACF-O5	Shadow AI	C2, C4
ACF-O6	Shadow agents	C5
ACF-O7	Shift of power	C1, C5
ACF-O8	Exceptions register	C3, C4
ACF-O9	Fleet governance	C5
ACF-O10	Affected people and redress	C1, C3
ACF-P1	Configuration grid	All
ACF-P2	Acceptance protocol	C1, C4, C5
ACF-P3	Ramp-up plan	C1, C4, C5
ACF-P4	Supervision log	C2, C3, C4, C5
ACF-P5	Role allocation	C2, C3, C5
ACF-P6	Escalation & trust	C3, C5
ACF-P7	Stop test	C1, C4, C5
ACF-P8	Incident log	C1, C4
ACF-P9	Permissions review	C3, C5
ACF-P10	Vendor assessment	C1, C5
ACF-P11	Foundation dependency	C4, C5
ACF-P12	Degraded mode & exit	C1, C3, C5
ACF-P13	Model version log	C1, C4
ACF-P14	Holding the record	C3, C5
ACF-P15	Blind spot statement	C4, C5



5 DELIVERABLE ASSESSMENT RUBRIC · OUT OF 100 POINTS

Scoring per criterion: **Excellent 4** · Good 3 · Satisfactory 2 · Insufficient 1 · weighted score scaled to 100.

CRITERION	EXCELLENT (4)	GOOD (3)	SATISFACTORY (2)	INSUFFICIENT (1)	WEIGHT
Understanding of the context	Comprehensive, relevant analysis	Strong analysis, minor omissions	Sound but incomplete analysis	Superficial analysis	20%
Use of ACF tools	Tools mastered and well applied	Tools used correctly	Tools poorly used	Tools ignored	20%
Quality of decisions & rationale	Decisions justified and aligned	Coherent decisions	Weakly argued decisions	Unjustified decisions	25%
Risk management & guardrails	Risks anticipated, robust guardrails	Risks identified and addressed	Risks only partially addressed	Risks ignored	20%
Presentation & teamwork	Clear, structured delivery	Adequate presentation	Confused presentation	No delivery	15%
Total score					/100

6 GRADING MATERIALS

EACH CASE IS PAIRED WITH A MODEL ANSWER SPECIFYING

- ✓ The **key expected elements**
- ✓ The **recommended decisions**
- ✓ The **standard justifications**
- ✓ The **common mistakes** to avoid
- ✓ The **best practices** to reward

INSTRUCTOR TIP

- Several answers may be valid as long as they are **coherent and justified** within the framework.
- Grading remains a **pedagogical act**: it is discussed with the students, never automated.



7 FACILITATE & ADAPT THE WORKSHOP

FACILITATION BEST PRACTICES

- ✓ Create a climate of trust and encourage everyone to participate.
- ✓ Ask open-ended questions to stimulate critical thinking.
- ✓ Challenge the groups with dilemmas that have no single right answer.
- ✓ Reward **justification** as much as the decision itself.
- ✓ Tie every choice back to the framework's **sovereignty** principles.

POSSIBLE ADAPTATIONS

- ✓ Cut the duration to 90 min by selecting fewer tools.
- ✓ Adapt the case's complexity level to the audience.
- ✓ Use a "flipped classroom" model (pre-work).
- ✓ Run a contest or a teaching hackathon.

8 RESOURCES FOR THE INSTRUCTOR

- ✓ ACF® presentation (slides).
- ✓ ACF® tool cards to hand out.
- ✓ Teaching cases (full scenarios).
- ✓ Detailed model answers & response elements.
- ✓ Deliverable templates (PDF / spreadsheet).
- ✓ Student card ACF-18 (guided practical case).

9 SAMPLE SESSION PLAN · CASE C1, E-COMMERCE PRICING AGENT

INTRODUCTION	CASE ANALYSIS	APPLICATION	DECISIONS	PRESENTATION	DEBRIEF
15 min	30 min	60 min	20 min	15 min	10 min
Introduce the case and the objectives · recap key ACF® principles.	Read the scenario · identify stakes and constraints.	Complete ACF-02, ACF-03, ACF-05, ACF-06, ACF-07.	Choose the autonomy level · limits and guardrails.	Delivery by each group.	Questions & discussion · key lessons to retain.



10 EXPECTED FINAL DELIVERABLE

PAGE 6/7

The deliverable produced by each group must cover the following seven items.

DELIVERABLE	CONTENT	DONE
Scenario analysis	Stakes, opportunities and constraints identified.	☐
Decision mapping	Decisions listed, assessed and ranked.	☐
Maturity and autonomy	Maturity of the agent as a whole (N0 → N3) and delegated autonomy decision by decision (suggestive → autonomous), each one justified.	☐
Guardrails & constraints	Human and technical guardrails defined.	☐
Indicators & supervision	KPIs and supervision cadence defined.	☐
Integration plan	5-phase plan with key actions.	☐
Summary & delivery	Clear summary of the approach and choices.	☐

FINAL TIP

- ✓ A successful rollout rests on the **balance between the agent's autonomy and the company's sovereignty**. Document, measure, control, adjust.

Use together with: **ACF-18** (guided practical case, to be distributed to students), **ACF-20** (governance simulation) and the full toolkit, **45 cards** in three families: Decide (ACF-O1 to ACF-O20), Organise (ACF-O1 to ACF-O10) and Operate (ACF-P1 to ACF-P15). The map in section 11 says which one is worked at which moment.

"Teaching agentic governance means training tomorrow's decision-makers."

ACF TOOLKIT® · INSTRUCTOR PACK



11 THE 45 CARDS AND WHEN EACH ONE IS WORKED

A workshop never uses all 45 cards. This map says which one serves which moment, so the instructor builds a session instead of handing out a stack.

Decide · ACF-01 to ACF-20

What is delegated and within which limits. This is the first-workshop family: without it there is nothing to organise or operate.

01 to 05	Diagnose, map, rate, frame, describe the agent
06 to 08	Supervise, stop, build the first agent briefing
09 to 12	Record, correct, audit, rate the risk
13 to 17	Mandate, anchor accountability, aim, test, review
18 to 20	The session documents: guided case, instructor's guide, simulation

Organise · ACF-O1 to ACF-O10

What the agent shifts inside the organisation. A second-level workshop, for participants who run a team.

01 to 04	Split human, AI and agent; map the hybrid; keep the critical skills
05 to 07	Write the AI policy; surface undeclared usage and undeclared agents
08 to 09	Record the shift of power; keep the exceptions register
O10	Govern a fleet; open redress to the people the agent affects

Operate · ACF-P1 to ACF-P15

What happens once the agent is live. A technical workshop, for participants who run a system.

P1 to P4	Configure, accept, ramp up, keep the supervision log
P5 to P8	Split the roles, escalate, test the stop, keep the incident log
P9 to P12	Review permissions, assess the vendor, measure dependency, prepare the exit
P13 to P15	Log the versions, hold the record, state the blind spots

Building a session: one case from table 3, the matching Decide tools, then a single Organise or Operate card chosen for the debate it opens. The toolkit's two running cases, Lendari and Rhena, each give a complete worked answer card by card: comparing two answers to the same question beats one model answer.



OBJECTIVE

Place participants in a situation of **agentic governance** facing critical events.
Develop the skills of **decision-making, arbitration, escalation and supervision** within a secure framework, and strengthen alignment with the sovereignty requirements of the ACF® framework.

1 PRINCIPLES OF THE SIMULATION

- ✓ **Realistic** : based on concrete cases drawn from the field.
- ✓ **Collaborative** : each role has its own mandate and objectives.
- ✓ **Adversarial** : tensions between business, risk and compliance.
- ✓ **Traceable** : all decisions are documented.
- ✓ **Aligned** : with the ACF® framework: autonomy levels, safeguards, escalations and duty of supervision.

2 GOVERNANCE ROLES

ROLE	CORE MISSION
DDAO (Delegated Decision Agent Officer)	Supervise, arbitrate and validate delegated decisions. Guarantor of controlled autonomy.
AI Governance Lead	Ensure the consistency of the ACF® framework within the organization and lead governance reviews.
Risk Officer	Identify, assess and monitor risks related to agents and decisions; recommend mitigations.
Compliance Officer	Ensure regulatory (AI Act, GDPR) and sectoral compliance, and the traceability of decisions.
Agent Owner (business)	Define business objectives, validate performance and ensure value generation.
Incident Manager	Manage incidents, crises and critical escalations; coordinate response and return to normal operations.

3 ACF DECISION CYCLE IN REAL CONDITIONS





4 THE CENTRAL ROLE OF THE DDAO

The DDAO, sovereign arbiter of the simulation

- ✓ Sovereign arbiter of **critical decisions** and autonomy levels.
- ✓ Guarantor of compliance with the ACF® framework and its safeguards.
- ✓ May escalate, demand evidence, **reduce autonomy or block** a decision.
- ✓ Responsible for overall supervision and the **protection of decisional sovereignty**.

5 ACF® AUTONOMY LEVELS · REFERENCE SCALE APPLIED TO EACH DECISION

SUGGESTIVE	CO-DECISION	SUPERVISED AUTONOMOUS	AUTONOMOUS
The agent proposes, the human decides.	Human validation before execution.	Acts alone, ex-post review.	Acts alone, without systematic review.

Autonomy is set **decision by decision** (ACF-02 / ACF-03). To be distinguished from the **overall maturity** of the agent (N0 → N3, ACF-05).

6 RUN SHEET FOR A STANDARD SIMULATION - 3-HOUR EXAMPLE

00:15 BRIEFING	00:35 INITIAL SCENARIO	01:15 EVENTS & ANALYSES	01:45 ARBITRATIONS	02:15 ACTIONS & MEASURES	03:00 DEBRIEF & LESSONS LEARNED
Presentation of the case, objectives, roles and rules.	Starting context and first elements to analyze.	Injection of incidents and signals; analysis by each role.	Proposals, tensions, arbitrations and governance choices.	Implementation of decisions and operational management.	Collective analysis, lessons learned and action plan.

Timings may be adjusted based on the complexity of the scenario and the maturity of the participants.

7 EXAMPLE SCENARIO · MARKETING CAMPAIGNS WITH DYNAMIC BUDGET

CONTEXT & OBJECTIVE

- A company uses an **autonomous agent** to run multi-channel marketing campaigns with a dynamic budget.
- Achieve performance targets while respecting budget, ethical and regulatory constraints.

CONSTRAINTS & STAKES

- Compliance with overall budget · AI Act and GDPR compliance · limited risk tolerance.
- Mandatory human supervision on certain actions.
- **Tensions:** growth vs. risk · performance vs. compliance · autonomy vs. control.



8 INJECTED EVENTS · EXAMPLES

- ◆ **Operational incident** - budget overrun detected on a campaign.
- ◆ **Compliance risk** - sensitive data collection without explicit consent.
- ◆ **Weak signal** - abnormal performance drop on a key channel.
- ◆ **Conflict of objectives** - business ROI in tension with ethical rules.
- ◆ **Crisis** - negative reaction on social media linked to a campaign.

9 CRITICALITY / AUTONOMY MATRIX

Criticality ▼ Autonomy ►	LOW CAPABILITY	MEDIUM CAPABILITY	HIGH CAPABILITY
HIGH CRITICALITY	CO-DECISION	SUGGESTIVE	SUGGESTIVE
MEDIUM CRITICALITY	SUPERVISED AUTONOMOUS	CO-DECISION	CO-DECISION
LOW CRITICALITY	SUPERVISED AUTONOMOUS	SUPERVISED AUTONOMOUS	AUTONOMOUS

SUGGESTIVE

 the human decides

CO-DECISION

 validation before execution

SUPERVISED AUTONOMOUS

 ex-post review

AUTONOMOUS

 no systematic review

The higher the criticality, the more delegated autonomy is reduced. Any critical decision remains arbitrated by the DDAO.

10 ESCALATION RULES & SUPERVISION CHECKPOINTS

QUICK ESCALATION RULES

- ✓ If **risk > tolerance** → automatic escalation to the DDAO.
- ✓ If **critical non-compliance** → immediate block.
- ✓ If **conflict of objectives** → human decision required.
- ✓ If **doubt about data** → require evidence or justification.
- ✓ If **reputational impact** → activate the crisis plan.

SUPERVISION CHECKPOINTS

- ✓ Before critical action.
- ✓ After major event.
- ✓ Before strategy change.
- ✓ After sensitive arbitration.
- ✓ Before cycle closure.



11 ⌚ SAFEGUARDS & NON-DELEGABLE DECISIONS

KILL SWITCH

- ✗ The DDAO may **reduce autonomy** or trigger a **Kill Switch (ACF-07)** if necessary. Any agent can be quarantined or deactivated immediately.

NON-DELEGABLE DECISIONS · EXAMPLES

- | | |
|--|--|
| ✓ Financial commitment beyond the approved threshold. | ✓ External communication that may impact reputation. |
| ✓ Collection / processing of sensitive data without consent. | ✓ Potential breach of law or regulations. |
| ✓ Modification of ethical rules or risk policies. | ✓ Full shutdown of a critical system or agent. |

12 📊 GOVERNANCE SCORECARD · OUT OF 600 POINTS

CRITERION ASSESSED	WHAT IS MEASURED	SCORE
Quality of decisions	Relevance & justification	/100
Risk management	Anticipation & mitigation	/100
Compliance with the ACF® framework	Autonomy, safeguards	/100
Responsiveness & relevance	Speed & accuracy	/100
Traceability & documentation	Decisions logged	/100
Collaboration & communication	Quality of arbitrations	/100
Overall governance score		/600

SCALE PER CRITERION · OUT OF 100

- | | |
|--|---|
| ● 0-49 Deficient · safeguard absent or ineffective | ● 50-74 Adequate · present but improvable |
| ● 75-100 Mastered · robust & documented | |

OVERALL SCORE · OUT OF 600

- | | | | |
|----------------------|-------------------------------|------------------------|---------------------|
| ● 0-199 Insufficient | ● 200-349 Needs strengthening | ● 350-499 Satisfactory | ● 500-600 Excellent |
|----------------------|-------------------------------|------------------------|---------------------|



13

KEY INDICATORS (KPIS) · OUT OF 600 POINTS

INDICATOR	DEFINITION	SCORE
ACF®-compliant decisions	% of decisions aligned with the framework.	/100
Average detection time	Time between event and detection.	/100
Appropriate escalation rate	Relevant escalations / total escalations.	/100
Effective corrective actions	Actions that resolved the issue.	/100
Autonomy level mastered	Autonomy used without critical incident.	/100
Quality of documentation	Decisions logged & justified.	/100
Overall KPI score		/600

14

ANALYSIS OF CRITICAL POINTS

CRITICAL ERRORS

Decisions or omissions that had a major impact on objectives or compliance.

POINTS OF VIGILANCE

Improvable decisions; risks controlled with room for improvement.

BEST PRACTICES TO CAPITALIZE ON

Relevant, responsive and compliant decisions · to be reused.



15

COMPLIANCE & SAFEGUARDS

CONTROL POINT	YES	PARTIAL	NO
Compliance with the ACF® framework	☐	☐	☐
AI Act & GDPR compliance	☐	☐	☐
Effective human supervision	☐	☐	☐
Documentation & traceability	☐	☐	☐
Adequate risk management	☐	☐	☐
Activation of safeguards	☐	☐	☐
Kill Switch used when necessary	☐	☐	☐

16

KEY LESSONS LEARNED

What worked well

What we will do differently

What needs improvement

Lessons for the organization

17

PRIORITY ACTION PLAN

NO.	PRIORITY ACTION	OWNER	DEADLINE
1			
2			
3			



18



AGENTIC GOVERNANCE MATURITY · CHECK THE LEVEL REACHED

1

INITIAL

Reactive, lightly formalized governance. Limited autonomy.



2

DEFINED

Roles and processes defined. Safeguards activated. Active supervision.



3

MASTERED

Autonomy adapted. Decisions logged. Effective escalations.



4

OPTIMIZED

Proactive governance. Indicators steered. Continuous improvement.



5

EXEMPLARY

Robust agentic sovereignty. Trust, resilience and compliance.



To be used with: ACF-18 (guided case study) and ACF-19 (instructor's guide). Tools mobilized: ACF-02/03 (mapping & criticality), ACF-05 (maturity), ACF-06 (supervision), ACF-07 (kill switch), ACF-12 (risks), ACF-13 (mandate). **Going further, by role:** the DDAO holds ACF-08 (exceptions register) and ACF-09 (fleet governance); the business holds ACF-P5 (role split) and ACF-P6 (escalation); engineering holds ACF-P7 (stop test), ACF-P12 (degraded mode and exit) and ACF-P14 (holding the record).

"You only govern well what you have learned to arbitrate under pressure."

ACF TOOLKIT® · GOVERNANCE SIMULATION

An agent never exists in a vacuum. This family looks at what it shifts around it: roles, the skills people stop practising, the usage nobody declared, and the power that changes hands without the org chart moving.

THE QUESTION THIS FAMILY ASKS

Who actually decides, now?

THE 10 CARDS

ACF-O1 Human–AI–Agent Matrix

ACF-O6 Shadow Agents Undeclared Agents

ACF-O2 Hybrid Organisation Map

ACF-O7 The Shift of Power

ACF-O3 Critical Skills to Keep

ACF-O8 The Exceptions Register

ACF-O4 AI Policy

ACF-O9 Governing the Fleet

ACF-O5 Shadow AI Undeclared Usage

ACF-O10 Affected People and Their Appeal

WHAT HAPPENS WHEN THIS FAMILY IS MISSING

- Agents run inside the organisation that governance never recorded (**ACF-O6**).
- A skill leaves the house because nobody practises it any more, and you notice only when you go looking for it (**ACF-O3**).
- A decision changes hands without the org chart moving: the named owner no longer decides anything (**ACF-O7**).

WHERE TO OPEN IT

ACF-O1 Who decides what, between the human, the rule and the agent.

ACF-O5 The AI usage nobody declared.

ACF-O9 The fleet, from the second agent onward.

Power does not disappear when a decision is automated, it changes hands. These ten cards exist to find out into whose.



OBJECTIVE

For every act handed to an agent: who performs it, who is accountable for it, which rule bounds it, who stays in the loop. You delegate execution, never accountability.

1 THE RULE · WHAT THE MATRIX MAKES VISIBLE

NON-DELEGATION OF ACCOUNTABILITY

You delegate execution, never accountability. The agent may hold the "R" (it performs); it never holds the "A" (it does not answer for its acts). The "A" always lands on a named human.

THE "RULE" COLUMN

Is not an actor that decides: it is the bound that makes autonomy governable, namely a cap, a stop trigger, a prohibition. It comes from the mandate (ACF-13) and the stop mechanism (ACF-07).

2 TO BE FILLED IN

PROCESS ACT	PERFORMS (R)	ACCOUNTABLE (A)	BOUNDING RULE	IN THE LOOP · AUTONOMY

3 CHECK WHEN RE-READING

"ACCOUNTABLE" COLUMN

Only human names, from the first act to the last, even though the agent performs three of the five. Accountability did not follow execution.

LAST ACT

Cutting the agent off is never performed by the agent. An organisation that let it decide its own shutdown would have no stop mechanism, only a suggestion.

The "Accountable" column holds only human names, never the agent. That is the test of this card.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 THE RULE · WHAT THE MATRIX MAKES VISIBLE

NON-DELEGATION OF ACCOUNTABILITY

You delegate execution, never accountability. The agent may hold the "R" (it performs); it never holds the "A" (it does not answer for its acts). The "A" always lands on a named human.

THE "RULE" COLUMN

Is not an actor that decides: it is the bound that makes autonomy governable, namely a cap, a stop trigger, a prohibition. It comes from the mandate (ACF-13) and the stop mechanism (ACF-07).

2 WORKED EXAMPLE · LENDARI · AFTER-SALES & RETURNS (AGENT MARGAUX, N2)

PROCESS ACT	PERFORMS (R)	ACCOUNTABLE (A)	BOUNDING RULE	IN THE LOOP · AUTONOMY
Answer a tier-1 after-sales request (tracking, delay, stock)	Agent	Léa Fontaine	Sourced answers; escalate outside perimeter	After-sales supervisor, a-posteriori check · supervised autonomy
Refund / credit note < €100	Agent	Léa Fontaine	Cap €100/action, €500/day (ACF-13); logging ACF-09	Daily human review · supervised autonomy
Grant a goodwill gesture (discount)	Agent proposes, human approves	Léa Fontaine	Never change catalogue prices; approval before sending	After-sales lead approves each case · co-decision
Handle a dispute / an out-of-perimeter request	Human (after-sales team)	Léa Fontaine	Mandatory human handover; the agent does not decide	The agent routes as best it can · suggestive
Cut the agent off in an incident	DDAO (Camille Roussel)	Camille Roussel	Trigger: error rate above 2% over 24h (ACF-07)	Never performed by the agent on itself · non-delegable zone



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

3 CHECK WHEN RE-READING

"ACCOUNTABLE" COLUMN

Only human names, from the first act to the last, even though the agent performs three of the five. Accountability did not follow execution.

LAST ACT

Cutting the agent off is never performed by the agent. An organisation that let it decide its own shutdown would have no stop mechanism, only a suggestion.

The "Accountable" column holds only human names, never the agent. That is the test of this card.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

1 THE RULE · WHAT THE MATRIX MAKES VISIBLE

NON-DELEGATION OF ACCOUNTABILITY

You delegate execution, never accountability. The agent may hold the "R" (it performs); it never holds the "A" (it does not answer for its acts). The "A" always lands on a named human.

THE "RULE" COLUMN

Is not an actor that decides: it is the bound that makes autonomy governable, namely a cap, a stop trigger, a prohibition. It comes from the mandate (ACF-13) and the stop mechanism (ACF-07).

2 WORKED EXAMPLE · RHENA · LEVEL 1 HOME CLAIMS (AGENT ANNA, N2)

PROCESS ACT	PERFORMS (R)	ACCOUNTABLE (A)	BOUNDING RULE	IN THE LOOP · AUTONOMY
Acknowledge the claim and chase missing documents	Agent	Ines Baumgartner	Two reminders at most, then a handler takes over	Review after the fact · supervised autonomy
Compute and propose settlement below €2,500	Agent	Ines Baumgartner	Cap €2,500/claim, €150,000/day (ACF-13); versioned depreciation schedule, logging ACF-09	Daily sampled review · supervised autonomy
Release the day-20 interim payment (§ 14 VVG)	Agent proposes, human approves	Marek Sobczak	Interim payment due at one month; never above the minimum owed	Handler approves each payment · joint decision
Decline cover, qualify a fraud	Human (home claims unit, integrity desk)	Marek Sobczak	The agent builds and reasons the file, it never rules	The agent prepares the file · advisory
Stop the agent during an incident	DDAO (Dr Katrin Ostermann)	Dr Katrin Ostermann	Schedule drift or error rate above 1.5% over 24h (ACF-07); authority valid across both entities	Never run by the agent on itself · non-delegable zone



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 CHECK WHEN RE-READING

"ACCOUNTABLE" COLUMN

Human names only, from the first act to the last, even though the agent performs two of the five on its own and prepares two more. Accountability did not follow execution.

LAST ACT

Stopping the agent is never performed by the agent. At Rhena the row carries one more clause, stop authority holds across both entities: on 11 March 2027, for want of that clause, Europe stopped in 3 minutes and the US subsidiary 40 minutes later.

The "Accountable" column holds only human names, never the agent. That is the test of this card.



OBJECTIVE

A one-page overview of how humans, agents and rules share the organisation, and of the domains where autonomy has settled in without governance. It locates; it does not score.

1 WHAT IT SHOWS · AND THE "NEEDS FRAMING" STATE

THE BLIND SPOT

Where autonomy has already slipped in unframed: a scoring engine, an allocation rule, an inherited automation nobody watches as an agent.

"NEEDS FRAMING" STATE

The most important on the map: here a decision is already being made without governance. A mature organisation does not aim for all-"Governed"; it aims for zero "needs framing" it ignores.

2 TO BE FILLED IN

DOMAIN / PROCESS	WHO PERFORMS	UNDER WHICH RULE	ACCOUNTABLE	STATE

3 WHAT IT REVEALS

READING YOUR ORGANISATION

One thing truly governed (the agent); two domains already deciding without a frame (recommendation, pricing); one clear boundary (payment above the cap is not delegated); the rest human.

NEXT TROUBLE

Not the agent you watch, already held, but the two "needs framing" cells nobody looks at as an agent. Each one lands in the action plan (ACF-10).

The real risk is not the agent you watch, but the "needs framing" domain nobody looks at as an agent.



RUNNING EXAMPLE

Lendari has a single governed agent so far, Margaux (after-sales, N2); the rest of the organisation is still human or algorithmic.

PAGE 1/1

1 WHAT IT SHOWS · AND THE "NEEDS FRAMING" STATE

THE BLIND SPOT

Where autonomy has already slipped in unframed: a scoring engine, an allocation rule, an inherited automation nobody watches as an agent.

"NEEDS FRAMING" STATE

The most important on the map: here a decision is already being made without governance. A mature organisation does not aim for all-"Governed"; it aims for zero "needs framing" it ignores.

2 WORKED EXAMPLE · LENDARI · PORTFOLIO VIEW

DOMAIN / PROCESS	WHO PERFORMS	UNDER WHICH RULE	ACCOUNTABLE	STATE
Customer relations · after-sales & returns	Agent "Margaux" (N2)	Supervised autonomy, cap €100/€500 (ACF-13)	Léa Fontaine	Governed
Product recommendation & merchandising	Algorithm, supervised	Decides in effect, no written mandate or ACF frame	E-commerce management	Needs framing
Pricing	Algorithm + human arbitration	Watched and switch-off-able, outside formal governance	Commercial management	Needs framing (sensitive)
Payment & refunds above the cap	Human	Non-delegable above €100/€500	Finance	Non-delegable
Marketing & content · logistics · HR	Human	—	Business leads	Human

3 WHAT IT REVEALS

READING LENDARI

One thing truly governed (Margaux); two domains already deciding without a frame (recommendation, pricing); one clear boundary (payment above the cap is not delegated); the rest human.

NEXT TROUBLE

Not the agent you watch, already held, but the two "needs framing" cells nobody looks at as an agent. Each one lands in the action plan (ACF-10).

The real risk is not the agent you watch, but the "needs framing" domain nobody looks at as an agent.



RUNNING EXAMPLE

Rhena draws its map in November 2026, before the first mandate. The depreciation schedule is flagged "to be framed", and nobody picks it up before the March 2027 incident.

PAGE 1/2

1 WHAT IT SHOWS · AND THE "NEEDS FRAMING" STATE

THE BLIND SPOT

Where autonomy has already slipped in unframed: a scoring engine, an allocation rule, an inherited automation nobody watches as an agent.

"NEEDS FRAMING" STATE

The most important on the map: here a decision is already being made without governance. A mature organisation does not aim for all-"Governed"; it aims for zero "needs framing" it ignores.

2 WORKED EXAMPLE · RHENA · PORTFOLIO VIEW, NOVEMBER 2026

DOMAIN / PROCESS	WHO PERFORMS	UNDER WHICH RULE	ACCOUNTABLE	STATE
Level 1 home claims, Europe	Agent "ANNA" (N2)	Supervised autonomy, cap €2,500/€150,000 (ACF-13)	Ines Baumgartner	Governed
Level 1 home claims, US subsidiary	The same agent, separate platform	No local mandate, the group policy does not bind Delaware	Brian McDouglas	To be framed (sensitive)
Fraud scoring at first notice of loss	Scoring model	Ranks and prioritises in fact, with no written mandate	Integrity desk	To be framed
Depreciation schedule and tariff revision	Actuarial chain	Revised once a year, outside agentic governance	Technical division	To be framed (sensitive)
Declining cover, qualifying a fraud	Human	Non-delegable, whatever the amount	Marek Sobczak	Non-delegable
Underwriting · reinsurance · investments	Human	—	Line managers	Human



RUNNING EXAMPLE

Rhena draws its map in November 2026, before the first mandate. The depreciation schedule is flagged "to be framed", and nobody picks it up before the March 2027 incident.

PAGE 2/2

3 WHAT IT REVEALS

RHENA READING

A single domain properly governed, and on one side of the Atlantic only; two chains already deciding unframed, the fraud score and the depreciation schedule; one clean boundary, declining cover is never delegated. The schedule row is flagged "to be framed" four months before it costs €77,400.

NEXT TROUBLE

Not the agent you watch, already held, but the two "needs framing" cells nobody looks at as an agent. Each one lands in the action plan (ACF-10).

The real risk is not the agent you watch, but the "needs framing" domain nobody looks at as an agent.



OBJECTIVE

The list of skills to keep in-house to stay able to supervise, audit and take back control of your agents. Sovereignty framing, not an HR training plan.

1 THE RULE THIS CARD PROTECTS

DELEGATE WITHOUT UNLEARNING

You can delegate a task without losing the skill to judge it. The day you no longer know how, you can no longer supervise. Stop doing long enough and the control gap becomes a competence gap.

2 TO BE FILLED IN

CRITICAL SKILL	WHY CRITICAL	LEVEL TO KEEP	DE-SKILLING SIGNAL

3 READING PRINCIPLE

CHOOSING A LEVEL IS GOVERNING

Not everything is kept at full mastery. Choosing a level per skill is already governing. What you do not choose, you lose by default.

WHAT THE CARD PROTECTS

Not the agent's performance, but the team's ability to judge it. As long as a human can handle a dispute and read an answer critically, the agent stays governable.

The most dangerous signal is never the agent's error: it is the control signed off without being read.



RUNNING EXAMPLE

Delegated domain at Lendari: after-sales & returns, handed to agent Margaux (N2).
Card kept by Camille Roussel (DDAO) with Léa Fontaine (after-sales).

PAGE 1/1

1 THE RULE THIS CARD PROTECTS

DELEGATE WITHOUT UNLEARNING

You can delegate a task without losing the skill to judge it. The day you no longer know how, you can no longer supervise. Stop doing long enough and the control gap becomes a competence gap.

2 WORKED EXAMPLE · LENDARI · AFTER-SALES & RETURNS (MARGAUX, N2)

CRITICAL SKILL	WHY CRITICAL	LEVEL TO KEEP	DE-SKILLING SIGNAL
Handle a complex after-sales dispute	Take back control (human handover)	Full mastery	No dispute handled by a human for months
Judge the quality of an after-sales answer	Supervise Margaux	Control capability	The supervisor signs off without reading; real control drops
Know the returns policy and its exceptions	Audit and explain a decision	Full mastery	Only the agent still knows the rules; no human referent left
Decide a goodwill gesture	Take back control (co-decision)	Full mastery	Moot as long as the decision stays a co-decision

3 READING PRINCIPLE

CHOOSING A LEVEL IS GOVERNING

Not everything is kept at full mastery. Choosing a level per skill is already governing. What you do not choose, you lose by default.

WHAT THE CARD PROTECTS

Not Margaux's performance, but the team's ability to judge it. As long as a human can handle a dispute and read an answer critically, the agent stays governable.

The most dangerous signal is never the agent's error: it is the control signed off without being read.



RUNNING EXAMPLE

Delegated domain at Rhena: first-level home claims, handed to ANNA (N2). Card held by Dr Katrin Ostermann (DDAO) with Ines Baumgartner.

PAGE 1/1

1 THE RULE THIS CARD PROTECTS

DELEGATE WITHOUT UNLEARNING

You can delegate a task without losing the skill to judge it. The day you no longer know how, you can no longer supervise. Stop doing long enough and the control gap becomes a competence gap.

2 WORKED EXAMPLE · RHENA · LEVEL 1 HOME CLAIMS (ANNA, N2)

CRITICAL SKILL	WHY CRITICAL	LEVEL TO KEEP	DE-SKILLING SIGNAL
Read a depreciation schedule and redo the sum by hand	Check what the agent proposes	Full command	Nobody can say any more where the applied rate comes from
Handle a disputed claim end to end	Take back control (human handover)	Full command	No disputed file handled by a human for months
Reason a declined cover	Non-delegable zone	Full command	Refusals copy out the note the agent prepared
Spot a fraud signal without the score	Supervise the integrity desk	Capacity to check	A file the score ranked low is never reopened
Defend a file before the Ombudsmann	Its ruling binds Rhena up to €10,000	Full command	Nobody can reconstruct the reasoning behind a decision

3 READING PRINCIPLE

CHOOSING A LEVEL IS GOVERNING

Not everything is kept at full mastery. Choosing a level per skill is already governing. What you do not choose, you lose by default.

WHAT THE CARD PROTECTS

Not ANNA's performance, but the team's capacity to judge it. The day Rhena has to explain before the Ombudsmann why a policyholder received €1,920 and not €2,100, a human will speak, and that human will have to be able to redo the sum.

The most dangerous signal is never the agent's error: it is the control signed off without being read.



OBJECTIVE

Give the organisation a single document covering two things that generic policies rarely separate: the daily use of AI tools, and the moment a system decides in place of a human. It is signed, it is revised, it comes before the first agent.

1 THE RULE · TWO REGIMES, ONE TEXT

USE AND DECISION

Assisted AI proposes and the human decides. Agentic AI decides and acts within a perimeter. Legal accountability sits in the second regime, which almost no usage policy addresses. The ACF policy covers both and defers to the agentic constitution (ACF-04) for the detail of each agent.

WHAT THE POLICY COMES BEFORE

It is laid down before the first agent and grounds the mandates (ACF-13) and the non-delegable zones. An organisation that writes mandates with no policy writes rules and nobody knows on whose authority they bind.

2 TO BE FILLED IN

PART	WHAT IT SETS	WHO OWNS IT	REVIEW

3 CHECK WHEN RE-READING

THE LIST OF PROHIBITIONS

It must name decisions, not technologies. A ban aimed at a product ages within six months; “no account closed without a human” holds.

THE REVIEW DATE

A policy with no review date written in becomes an archive document the day the tooling changes.

A policy that does not list the decisions it refuses to delegate is not a policy, it is an intention.



RUNNING EXAMPLE

Lendari writes its AI policy in March 2026, three months before Margaux's mandate.
Owned by Camille Roussel (DDAO).

PAGE 1/1

1 THE RULE • TWO REGIMES, ONE TEXT

USE AND DECISION

Assisted AI proposes and the human decides. Agentic AI decides and acts within a perimeter. Legal accountability sits in the second regime, which almost no usage policy addresses. The ACF policy covers both and defers to the agentic constitution (ACF-04) for the detail of each agent.

WHAT THE POLICY COMES BEFORE

It is laid down before the first agent and grounds the mandates (ACF-13) and the non-delegable zones. An organisation that writes mandates with no policy writes rules and nobody knows on whose authority they bind.

2 WORKED EXAMPLE • LENDARI • AI POLICY V1.0, MARCH 2026

PART	WHAT IT SETS	WHO OWNS IT	REVIEW
A • Tool use	Approved tools, data barred from prompts, review before any customer-facing send	IT, with compliance	Yearly, or on any new tool
A • Data	No named customer data in a tool not hosted in the EU	DPO	Yearly
B • Delegated decision	Every deciding agent goes through a written mandate (ACF-13) and a rated criticality (ACF-03)	Camille Roussel (DDAO)	On every new mandate
B • Non-delegable zones	Closing an account, reporting fraud, answering an open dispute	Customer director	Twice a year
C • Governance	Quarterly review, decision register (ACF-09), named route of appeal (ACF-O10)	Camille Roussel (DDAO)	Quarterly

3 CHECK WHEN RE-READING

THE LIST OF PROHIBITIONS

It must name decisions, not technologies. A ban aimed at a product ages within six months; “no account closed without a human” holds.

THE REVIEW DATE

A policy with no review date written in becomes an archive document the day the tooling changes. At Lendari the review is owned by a named person, not by a department.

A policy that does not list the decisions it refuses to delegate is not a policy, it is an intention.



RUNNING EXAMPLE

Rhena settles its AI policy in October 2026, three months before ANNA's mandate.
Owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE RULE · TWO REGIMES, ONE TEXT

USE AND DECISION

Assisted AI proposes and the human decides. Agentic AI decides and acts within a perimeter. Legal accountability sits in the second regime, which almost no usage policy addresses. The ACF policy covers both and defers to the agentic constitution (ACF-04) for the detail of each agent.

WHAT THE POLICY COMES BEFORE

It is laid down before the first agent and grounds the mandates (ACF-13) and the non-delegable zones. An organisation that writes mandates with no policy writes rules and nobody knows on whose authority they bind.

2 WORKED EXAMPLE · RHENA · AI POLICY V1.0, OCTOBER 2026

PART	WHAT IT SETS	WHO OWNS IT	REVIEW
A · Tool use	Approved tools, human read-through before anything reaches a policyholder	IT, with compliance	Yearly, or at any new tool
A · Data	No health data and no loss-adjuster report in a tool hosted outside the European Union	Group DPO	Yearly
B · Delegated decision	Every deciding agent goes through a written mandate (ACF-13) and a rated criticality (ACF-03)	Dr Katrin Ostermann (DDAO)	At every new mandate
B · Non-delegable zones	Declining cover, qualifying a fraud, files in litigation, policyholders under guardianship	Marek Sobczak	Twice a year
B · Territorial reach	The group policy is not a mandate outside Germany: each entity signs its own	Legal division	At every new country
C · Governance	Quarterly review, decision register (ACF-09), named route of appeal (ACF-O10)	Dr Katrin Ostermann (DDAO)	Quarterly



RUNNING EXAMPLE

Rhena settles its AI policy in October 2026, three months before ANNA's mandate.
Owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 CHECK WHEN RE-READING

THE LIST OF PROHIBITIONS

It must name decisions, not technologies. A ban aimed at a product ages within six months; “no account closed without a human” holds.

THE REVIEW DATE

A policy with no review date written into it becomes an archive document the day the tooling changes. At Rhena the review is owned by a named person rather than a department, and the territorial reach row was added afterwards, when the US subsidiary asked for the same agent.

A policy that does not list the decisions it refuses to delegate is not a policy, it is an intention.



OBJECTIVE

Flush out the AI tools used outside any framework, sort them, and decide for each whether it enters the policy (ACF-O4) or stops. This card covers tools that assist a person, not agents that decide: those belong to the next card.

1 THE RULE · LOOK FOR THE USAGE, NOT THE CULPRIT

WHY USAGE HIDES

Someone who saves 2 hours a day does not give it up because a memo forbids it, they stop mentioning it. The sweep therefore runs with no sanction announced, otherwise it measures people's caution rather than their practice.

WHAT IS MEASURED

Three things only: which data leaves the organisation, which act reaches a third party unreviewed, and which work already depends on a tool nobody maintains.

2 TO BE FILLED IN

USAGE SURFACED	WHO	DATA EXPOSED	DECISION

3 CHECK WHEN RE-READING

THE FIGURE THAT MATTERS

Not the number of tools found, but the number of uses stopped with nothing put in their place.

THE FREQUENCY

The sweep is repeated, because usage re-forms. A single sweep gives a snapshot; two give a trend, which is the only useful thing.

Banning without offering an alternative does not remove the usage, it moves it out of sight.



RUNNING EXAMPLE

Lendari sweeps its undeclared usage in April 2026, before deploying Margaux. 47 uses surfaced in 3 weeks.

PAGE 1/1

1 THE RULE · LOOK FOR THE USAGE, NOT THE CULPRIT

WHY USAGE HIDES

Someone who saves 2 hours a day does not give it up because a memo forbids it, they stop mentioning it. The sweep therefore runs with no sanction announced, otherwise it measures people's caution rather than their practice.

WHAT IS MEASURED

Three things only: which data leaves the organisation, which act reaches a third party unreviewed, and which work already depends on a tool nobody maintains.

2 WORKED EXAMPLE · LENDARI · APRIL 2026 SWEEP

USAGE SURFACED	WHO	DATA EXPOSED	DECISION
Drafting support replies in a consumer assistant	6 people in support	Customer name, order and address	Stopped, replaced by the approved tool
Translating product sheets	Content team	No personal data	Kept, entered in the register
Analysing customer reviews	Marketing	Public reviews, pseudonyms	Kept, on condition of anonymised export
Generating campaign visuals	Studio	None	Kept, disclosure required on publication
Screening job applications	1 recruiter	CVs, sensitive data	Stopped, non-delegable zone (ACF-O4, part B)

3 CHECK WHEN RE-READING

THE FIGURE THAT MATTERS

Not the number of tools found, but the number of uses stopped with nothing put in their place. At Lendari, 2 out of 47, and each got an approved tool within the month.

THE FREQUENCY

The sweep is repeated, because usage re-forms. A single sweep gives a snapshot; two give a trend, which is the only useful thing.

Banning without offering an alternative does not remove the usage, it moves it out of sight.



RUNNING EXAMPLE

Rhena sweeps its undeclared uses in September 2026, before ANNA's mandate. 214 uses surfaced in 5 weeks, across 28 countries.

1 THE RULE · LOOK FOR THE USAGE, NOT THE CULPRIT

WHY USAGE HIDES

Someone who saves 2 hours a day does not give it up because a memo forbids it, they stop mentioning it. The sweep therefore runs with no sanction announced, otherwise it measures people's caution rather than their practice.

WHAT IS MEASURED

Three things only: which data leaves the organisation, which act reaches a third party unreviewed, and which work already depends on a tool nobody maintains.

2 WORKED EXAMPLE · RHENA · SEPTEMBER 2026 SWEEP

USAGE SURFACED	WHO	DATA EXPOSED	DECISION
Drafting refusal letters in a consumer assistant	34 handlers, 6 countries	Name, address, amount and grounds for refusal	Stopped, replaced by the approved tool
Summarising loss-adjuster reports	Home claims unit	The full report, at times health material	Stopped, non-delegable zone (ACF-O4, part A)
Translating policy wordings	Legal team	No personal data	Kept, entered in the register
Analysing complaint verbatims	Quality	Pseudonymised verbatims	Kept, on condition of anonymised export
Screening job applications	2 recruiters, US subsidiary	CVs, sensitive data	Stopped, non-delegable zone

3 CHECK WHEN RE-READING

THE FIGURE THAT MATTERS

It is not the number of tools found, it is the number of uses stopped with nothing to replace them. At Rhena, 9 out of 214, and it took three months to replace them because an approved tool has to be approved in 28 countries and not in one.

THE FREQUENCY

The sweep is repeated, because usage re-forms. A single sweep gives a snapshot; two give a trend, which is the only useful thing.

Banning without offering an alternative does not remove the usage, it moves it out of sight.



OBJECTIVE

Find the agents acting in the organisation's name with no mandate: the ones an employee plugged in, the ones a supplier switched on through an update, the ones an open connector lets write. An agent without a mandate binds the organisation all the same.

1 THE RULE · WHAT TO LOOK FOR AND HOW IT GETS IN

TOOL OR AGENT

A tool proposes and stops. An agent acts: it sends, it changes, it commits. The question is not whether it is clever, but whether it produces an act the organisation will have to answer for before a third party.

THE THREE WAYS IN

The employee who plugs an assistant into their mailbox. The supplier who switches on an agentic feature through an update. The connector opened to a third party, which hands an outside agent a right to write.

2 TO BE FILLED IN

AGENT FOUND	HOW IT GOT IN	ACT IT PERFORMS	DECISION

3 CHECK WHEN RE-READING

THE MANDATE TEST

For each agent found, one question only: is there a named human who answers for what it does? If the answer is slow to come, the agent stops until the person is found.

WHAT FOLLOWS THE CENSUS

A kept agent is not a regularised agent until it has a mandate (ACF-13), a rated criticality (ACF-03) and a tested stop mechanism (ACF-07).

An undeclared agent becomes visible only at the first incident, and the incident always comes before the declaration.



RUNNING EXAMPLE

In May 2026 Lendari finds four agents plugged into its tools, none of which holds a mandate.

PAGE 1/1

1 THE RULE · WHAT TO LOOK FOR AND HOW IT GETS IN

TOOL OR AGENT

A tool proposes and stops. An agent acts: it sends, it changes, it commits. The question is not whether it is clever, but whether it produces an act the organisation will have to answer for before a third party.

THE THREE WAYS IN

The employee who plugs an assistant into their mailbox. The supplier who switches on an agentic feature through an update. The connector opened to a third party, which hands an outside agent a right to write.

2 WORKED EXAMPLE · LENDARI · MAY 2026 CENSUS

AGENT FOUND	HOW IT GOT IN	ACT IT PERFORMS	DECISION
A product manager's mail assistant	Extension installed by the person	Replies and files with no review	Cut off, write access withdrawn
Abandoned-basket follow-up	Switched on by a platform update	Sends the customer a discount	Mandate written (ACF-13), cap set
The carrier's returns-sorting agent	Connector opened to the provider	Approves or refuses a return	Suspended pending permissions review (ACF-P8)
Listing-writing assistant	Personal account plugged into the CRM	Publishes without approval	Cut off, usage moved to the approved tool

3 CHECK WHEN RE-READING

THE MANDATE TEST

For each agent found, one question only: is there a named human who answers for what it does? If the answer is slow to come, the agent stops until the person is found.

WHAT FOLLOWS THE CENSUS

A kept agent is not a regularised agent until it has a mandate (ACF-13), a rated criticality (ACF-03) and a tested stop mechanism (ACF-07).

An undeclared agent becomes visible only at the first incident, and the incident always comes before the declaration.



1 THE RULE · WHAT TO LOOK FOR AND HOW IT GETS IN

TOOL OR AGENT

A tool proposes and stops. An agent acts: it sends, it changes, it commits. The question is not whether it is clever, but whether it produces an act the organisation will have to answer for before a third party.

THE THREE WAYS IN

The employee who plugs an assistant into their mailbox. The supplier who switches on an agentic feature through an update. The connector opened to a third party, which hands an outside agent a right to write.

2 WORKED EXAMPLE · RHENA · NOVEMBER 2026 CENSUS

AGENT FOUND	HOW IT GOT IN	ACT IT PERFORMS	DECISION
A unit head's mailbox assistant	Extension installed by the person themselves	Answers the policyholder and files without review	Cut off, write access withdrawn
Chasing missing documents	Switched on by a platform update	Writes to the policyholder in Rhena's name	Mandate written (ACF-13), two reminders at most
Pre-sorting by the loss-adjuster network	Connector opened to the provider	Ranks the file before Rhena ever sees it	Suspended pending a permissions review (ACF-P8)
Bank reconciliation of settlements	Service account left over from a closed project	Posts settlement entries	Cut off, returned to four-eyes approval
Translating outgoing letters, US subsidiary	Personal account wired into the CRM	Sends with no approval	Cut off, use moved to the approved tool

3 CHECK WHEN RE-READING

THE MANDATE TEST

For each agent found, one question only: is there a named human who answers for what it does? If the answer is slow to come, the agent stops until the person is found.

WHAT FOLLOWS THE CENSUS

A kept agent is not a regularised agent until it has a mandate (ACF-13), a rated criticality (ACF-03) and a tested stop mechanism (ACF-07).

An undeclared agent becomes visible only at the first incident, and the incident always comes before the declaration.



OBJECTIVE

Make visible the power that has moved. Every agent deployed takes a decision away from someone and hands another to someone else, usually with that transfer neither discussed nor even noticed.

1 THE RULE · WHAT ACTUALLY MOVES

THE POWER TO CHALLENGE

Rarely the decision itself, almost always the power to challenge it. The person who used to decide keeps the title and loses the grip, while the person who configures the agent never received a mandate to decide and decides all the same.

A SETTING IS AN ACT

Setting a threshold decides thousands of cases in advance. Configuration is therefore a governance act and falls under the same rules as the decision it replaces.

2 TO BE FILLED IN

DECISION	WHO DECIDED BEFORE	WHO ACTUALLY DECIDES	WHAT MUST BE RE-ESTABLISHED

3 CHECK WHEN RE-READING

THE MIDDLE COLUMN

It is the only one that counts. If “who actually decides” holds neither a human nor a written rule, the organisation has delegated without knowing it.

WHAT THE CARD TRIGGERS

Every line where power moved with no written act ends in a mandate (ACF-13), a recorded exception (ACF-O8) or a reversal.

Power does not vanish when a decision is automated, it changes hands and stops being discussed.



RUNNING EXAMPLE

In July 2026 Lendari measures what six months of agents have changed in its decision paths.

1 THE RULE · WHAT ACTUALLY MOVES

THE POWER TO CHALLENGE

Rarely the decision itself, almost always the power to challenge it. The person who used to decide keeps the title and loses the grip, while the person who configures the agent never received a mandate to decide and decides all the same.

A SETTING IS AN ACT

Setting a threshold decides thousands of cases in advance. Configuration is therefore a governance act and falls under the same rules as the decision it replaces.

2 WORKED EXAMPLE · LENDARI · SIX MONTHS AFTER MARGAUX

DECISION	WHO DECIDED BEFORE	WHO ACTUALLY DECIDES	WHAT MUST BE RE-ESTABLISHED
Granting a refund under €100	Support adviser, case by case	The agent, within the cap	Signed mandate (ACF-13), cap revisited at every validity review
Setting the €100 cap	Nobody, the figure came with the tool	The integrator who configured it	Decision raised to the DDAO and entered in the register (ACF-09)
Escalating a dispute	Support adviser	The model's confidence threshold	Threshold documented, revisited at every model version
Refusing a goodwill gesture	Support manager	The agent, by silence, when it offers nothing	The implicit refusal becomes a recorded refusal
Stopping the agent	Not applicable	The DDAO alone	Power confirmed, deputy named in writing

3 CHECK WHEN RE-READING

THE MIDDLE COLUMN

It is the only one that counts. If “who actually decides” holds neither a human nor a written rule, the organisation has delegated without knowing it.

WHAT THE CARD TRIGGERS

Every line where power moved with no written act ends in a mandate (ACF-13), a recorded exception (ACF-O8) or a reversal.

Power does not vanish when a decision is automated, it changes hands and stops being discussed.



RUNNING EXAMPLE

In September 2027 Rhena measures what six months of an agent have changed in its decision paths.

PAGE 1/1

1 THE RULE • WHAT ACTUALLY MOVES

THE POWER TO CHALLENGE

Rarely the decision itself, almost always the power to challenge it. The person who used to decide keeps the title and loses the grip, while the person who configures the agent never received a mandate to decide and decides all the same.

A SETTING IS AN ACT

Setting a threshold decides thousands of cases in advance. Configuration is therefore a governance act and falls under the same rules as the decision it replaces.

2 WORKED EXAMPLE • RHENA • SIX MONTHS AFTER ANNA

DECISION	WHO DECIDED BEFORE	WHO ACTUALLY DECIDES	WHAT MUST BE RE-ESTABLISHED
Settle below €2,500	The handler, case by case	The agent, within the cap	Signed mandate (ACF-13), cap revisited at every validity review
Set the depreciation rate applied	The technical division, once a year	Whichever version of the schedule is loaded in the agent	Version entered in the register (ACF-09) and checked at every tariff revision
Decide that a file is complete	The handler	The agent's completeness check	Written rule, two reminders at most
Escalate a file	The handler	The model's confidence threshold	Documented threshold, revisited at every model version
Stop the agent	Not applicable	The DDAO for Europe, nobody for the subsidiary	Stop authority extended in writing to both entities

3 CHECK WHEN RE-READING

THE MIDDLE COLUMN

It is the only one that counts. If “who actually decides” holds neither a human nor a written rule, the organisation has delegated without knowing it.

WHAT THE CARD TRIGGERS

Every row where power moved with no written act ends in a mandate (ACF-13), a recorded exception (ACF-O8) or a step back. At Rhena the second row was written in September 2027, six months after a version of the schedule decided 430 settlements on its own.

Power does not vanish when a decision is automated, it changes hands and stops being discussed.



OBJECTIVE

Keep the list of departures granted from the rule: who asked, who granted, for how long, and on what condition they close. An exception with no end date becomes the rule, with nobody having had to vote it.

1 THE RULE • WHY WRITE THEM DOWN

SPOKEN AND RULE LOOK ALIKE

An exception granted verbally is indistinguishable from a rule. Six months later nobody knows whether the cap exceeded that day was a favour or the practice, and the audit can no longer tell.

WHAT NEVER TAKES AN EXCEPTION

A non-delegable zone admits no exception, otherwise it is not one. The register therefore separates what is negotiable, thresholds and delays, from what is not.

2 TO BE FILLED IN

EXCEPTION	REQUESTED BY	GRANTED BY	DEADLINE AND OUTCOME

3 CHECK WHEN RE-READING

THE DEADLINE COLUMN

No line stays empty. An exception with no date closes by default at the next review, and whoever wants it must ask again.

THE RECURRING EXCEPTION

Granted 3 times, an exception is no longer one: either the rule is wrong and gets rewritten, or the request is unfounded and gets refused for good.

An organisation does not drift through its rules, it drifts through the exceptions it forgot to close.



RUNNING EXAMPLE

Lendari has kept its exceptions register since June 2026. 9 open, 4 closed.

PAGE 1/1

1 THE RULE · WHY WRITE THEM DOWN

SPOKEN AND RULE LOOK ALIKE

An exception granted verbally is indistinguishable from a rule. Six months later nobody knows whether the cap exceeded that day was a favour or the practice, and the audit can no longer tell.

WHAT NEVER TAKES AN EXCEPTION

A non-delegable zone admits no exception, otherwise it is not one. The register therefore separates what is negotiable, thresholds and delays, from what is not.

2 WORKED EXAMPLE · LENDARI · REGISTER AT 09/30/2026

EXCEPTION	REQUESTED BY	GRANTED BY	DEADLINE AND OUTCOME
Cap raised to €250 during the sales	Sales director	Camille Roussel (DDAO)	Closed 02/15/2026, not renewed
Automatic Sunday replies with no supervision	Customer service	Léa Fontaine	Open, deadline 12/31/2026
CRM access for the carrier's agent	Logistics	IT, with no DDAO opinion	Irregular, taken back to permissions review (ACF-P8)
Dual control suspended during the 06/05 incident	Operations	Camille Roussel (DDAO)	Closed 06/12, reason written in the register
Acknowledgements sent without review	Customer service	Léa Fontaine	Open, revisited at every validity review (ACF-17)

3 CHECK WHEN RE-READING

THE DEADLINE COLUMN

No line stays empty. An exception with no date closes by default at the next review, and whoever wants it must ask again.

THE RECURRING EXCEPTION

Granted 3 times, an exception is no longer one: either the rule is wrong and gets rewritten, or the request is unfounded and gets refused for good.

An organisation does not drift through its rules, it drifts through the exceptions it forgot to close.



RUNNING EXAMPLE

Rhena has kept its exceptions register since February 2027. 14 open, 6 closed.

PAGE 1/1

1 THE RULE · WHY WRITE THEM DOWN

SPOKEN AND RULE LOOK ALIKE

An exception granted verbally is indistinguishable from a rule. Six months later nobody knows whether the cap exceeded that day was a favour or the practice, and the audit can no longer tell.

WHAT NEVER TAKES AN EXCEPTION

A non-delegable zone admits no exception, otherwise it is not one. The register therefore separates what is negotiable, thresholds and delays, from what is not.

2 WORKED EXAMPLE · RHENA · REGISTER AT 30/09/2027

EXCEPTION	REQUESTED BY	GRANTED BY	DEADLINE AND OUTCOME
Cap raised to €4,000 after the January storms	Claims division	Dr Katrin Ostermann (DDAO)	Closed 28/02/2027, not renewed
Daily review suspended at weekends	Home claims unit	Ines Baumgartner	Open, expires 31/12/2027
Full file access for the loss-adjuster network's agent	Operations	IT, without the DDAO's view	Irregular, pulled into a permissions review (ACF-P8)
Four-eyes check suspended during the 11/03 incident	Operations	Dr Katrin Ostermann (DDAO)	Closed 13/03, grounds written into the register
\$10,000 cap requested by the US subsidiary	Brian McDouglas	Refused	Reasoned refusal, revisited at the validity review (ACF-17)

3 CHECK WHEN RE-READING

THE DEADLINE COLUMN

No line stays empty. An exception with no date closes by default at the next review, and whoever wants it must ask again.

THE RECURRING EXCEPTION

Granted 3 times, an exception is no longer one: either the rule is wrong and gets rewritten, or the request is unfounded and gets refused for good.

An organisation does not drift through its rules, it drifts through the exceptions it forgot to close.



OBJECTIVE

Hold a population of agents rather than one agent. List what exists, spot the agents that talk to each other and the shared dependencies, and set the point beyond which the organisation can no longer say what it has deployed.

1 THE RULE · WHAT NUMBERS CHANGE

RISK BECOMES SHARED

A single agent has an owner and a stop mechanism. Ten agents share a platform, a supplier, a model version. Risk stops being individual and one outage reaches them all.

AGENTS THAT TALK TO EACH OTHER

Two agents calling each other form a chain whose mandate nobody wrote. One's output becomes the other's input, and each supervisor sees only half the path.

2 TO BE FILLED IN

AGENT	CRITICALITY	MANDATE AND OWNER	SHARED DEPENDENCY

3 CHECK WHEN RE-READING

THE RIGHT-HAND COLUMN

It says what falls together.

THE INVENTORY THRESHOLD

The day nobody can name every agent in service from memory, the inventory stops being a formality and becomes the only way to know what the organisation has delegated.

The first agent is governed by attention. The tenth is governed by inventory, or it is not governed at all.



RUNNING EXAMPLE

Lendari goes from 1 to 7 agents between June and December 2026, and finds its rules only held one at a time.

PAGE 1/1

1 THE RULE · WHAT NUMBERS CHANGE

RISK BECOMES SHARED

A single agent has an owner and a stop mechanism. Ten agents share a platform, a supplier, a model version. Risk stops being individual and one outage reaches them all.

AGENTS THAT TALK TO EACH OTHER

Two agents calling each other form a chain whose mandate nobody wrote. One's output becomes the other's input, and each supervisor sees only half the path.

2 WORKED EXAMPLE · LENDARI · FLEET AT 12/31/2026

AGENT	CRITICALITY	MANDATE AND OWNER	SHARED DEPENDENCY
Margaux, support and returns	High	ACF-13 of 06/15, Léa Fontaine	Model v1.2, shared platform
Abandoned-basket follow-up	Moderate	ACF-13 of 09/02, marketing	Model v1.2, customer base
Carrier returns sorting	High	None, being regularised (ACF-O6)	Provider's connector
Supplier restocking	Moderate	ACF-13 of 10/20, purchasing	Shared platform
Price watch	Low	ACF-13 of 10/20, marketing	External supplier

3 CHECK WHEN RE-READING

THE RIGHT-HAND COLUMN

It says what falls together. At Lendari 3 agents out of 5 depend on the same platform: one outage stops them all, and the degraded mode (ACF-P11) must be written once for the fleet.

THE INVENTORY THRESHOLD

The day nobody can name every agent in service from memory, the inventory stops being a formality and becomes the only way to know what the organisation has delegated.

The first agent is governed by attention. The tenth is governed by inventory, or it is not governed at all.



LEVELOrganisation

DURATION1 h 15 - 2 h 30

USEInventory and review

TRACKToolkit

RUNNING EXAMPLE

Rhena goes from 1 to 6 agents between January and December 2027, and finds its rules only held for a single one.

1 THE RULE · WHAT NUMBERS CHANGE

RISK BECOMES SHARED

A single agent has an owner and a stop mechanism. Ten agents share a platform, a supplier, a model version. Risk stops being individual and one outage reaches them all.

AGENTS THAT TALK TO EACH OTHER

Two agents calling each other form a chain whose mandate nobody wrote. One's output becomes the other's input, and each supervisor sees only half the path.

2 WORKED EXAMPLE · RHENA · FLEET AT 31/12/2027

AGENT	CRITICALITY	MANDATE AND OWNER	SHARED DEPENDENCY
ANNA, level 1 home claims Europe	High	ACF-13 of 15/01, Ines Baumgartner	Model v3.4, Frankfurt platform
ANNA, US subsidiary	High	Local mandate of 01/07, Brian McDouglas	Model v3.4, Virginia platform
Chasing missing documents	Moderate	ACF-13 of 02/05, home claims unit	Model v3.4, Frankfurt platform
Fraud scoring at first notice of loss	High	None, being regularised (ACF-O6)	Shared claims database
Assigning adjusters across the network	Moderate	ACF-13 of 20/09, technical division	Frankfurt platform
Translating outgoing letters	Low	ACF-13 of 20/09, communications	External provider

3 CHECK WHEN RE-READING

THE RIGHT-HAND COLUMN

It says what falls together. At Rhena, 4 of the 6 agents depend on the Frankfurt platform, and the degraded mode (ACF-P11) has to be written once for the fleet. It says something else too: the first two rows carry the same model on two platforms, so a fix applied to one leaves the other untouched.

THE INVENTORY THRESHOLD

The day nobody can name every agent in service from memory, the inventory stops being a formality and becomes the only way to know what the organisation has delegated.

The first agent is governed by attention. The tenth is governed by inventory, or it is not governed at all.



OBJECTIVE

Name who bears the delegated decisions, how that person learns a system decided, and before which human they can ask for a review. An appeal that leads only back to the same agent is not an appeal.

1 THE RULE · THREE QUESTIONS, ONE USEFUL APPEAL

THE THREE QUESTIONS

Who bears it. How that person learns of it. Before whom they can come back. Until the three answers are written, the organisation does not know what its delegation costs those who never agreed to it.

WHAT MAKES AN APPEAL USEFUL

It reaches a human with the power to undo the decision, within a written deadline, without the claimant having to prove the system at fault. A form that routes back to the same agent closes the appeal instead of opening it.

2 TO BE FILLED IN

DECISION	WHO BEARS IT	HOW THEY LEARN OF IT	APPEAL AND DEADLINE

3 CHECK WHEN RE-READING

THE LINE WITH NO APPEAL

Any line whose right-hand column is empty is a decision the organisation takes with no means of correcting it.

THE APPEAL NOBODY USES

An appeal never used is not a good sign: either nobody knows it exists, or the path is too long. The number of appeals filed reads as a health indicator, not as a cost.

A decision with no appeal is not an automated decision, it is a decision beyond appeal.



RUNNING EXAMPLE

In October 2026 Lendari opens a named route of appeal for Margaux's decisions.

PAGE 1/1

1 THE RULE · THREE QUESTIONS, ONE USEFUL APPEAL

THE THREE QUESTIONS

Who bears it. How that person learns of it. Before whom they can come back. Until the three answers are written, the organisation does not know what its delegation costs those who never agreed to it.

WHAT MAKES AN APPEAL USEFUL

It reaches a human with the power to undo the decision, within a written deadline, without the claimant having to prove the system at fault. A form that routes back to the same agent closes the appeal instead of opening it.

2 WORKED EXAMPLE · LENDARI · APPEALS IN OCTOBER 2026

DECISION	WHO BEARS IT	HOW THEY LEARN OF IT	APPEAL AND DEADLINE
Refund refused	The customer	A message naming the agent and the reason	Review by an adviser within 5 working days
Goodwill gesture refused by silence	The customer	They do not learn of it	To be fixed: refusal made explicit by 11/30
Block on suspicion of fraud	The customer	A generic message, reason withheld	Human review within 48 h, reason given
Return refused by the carrier's agent	The customer	A notification from the carrier	No appeal at Lendari, to be opened
Prioritising a case in the queue	The customer	Never	Outside appeal, but entered in the register (ACF-09)

3 CHECK WHEN RE-READING

THE LINE WITH NO APPEAL

Any line whose right-hand column is empty is a decision the organisation takes with no means of correcting it. At Lendari, 2 lines out of 5 in October, and that is the plan for the next two months.

THE APPEAL NOBODY USES

An appeal never used is not a good sign: either nobody knows it exists, or the path is too long. The number of appeals filed reads as a health indicator, not as a cost.

A decision with no appeal is not an automated decision, it is a decision beyond appeal.



RUNNING EXAMPLE

In April 2027 Rhena opens a named route of appeal against ANNA's decisions, one month after the schedule incident.

PAGE 1/1

1 THE RULE · THREE QUESTIONS, ONE USEFUL APPEAL

THE THREE QUESTIONS

Who bears it. How that person learns of it. Before whom they can come back. Until the three answers are written, the organisation does not know what its delegation costs those who never agreed to it.

WHAT MAKES AN APPEAL USEFUL

It reaches a human with the power to undo the decision, within a written deadline, without the claimant having to prove the system at fault. A form that routes back to the same agent closes the appeal instead of opening it.

2 WORKED EXAMPLE · RHENA · APPEALS IN APRIL 2027

DECISION	WHO BEARS IT	HOW THEY LEARN OF IT	APPEAL AND DEADLINE
Settlement computed below the expected amount	The policyholder	Letter naming the agent, the schedule and its version	Human re-examination within 15 days, then the Ombudsmann
File closed as incomplete	The policyholder	Generic message after two reminders	To be fixed: detailed grounds required by 30/06
Low priority set by the fraud score	The policyholder	They never learn of it	Human re-examination within 48h once they complain
File pre-sorted by the adjuster network	The policyholder	Notice from the adjuster	No appeal at Rhena, to be opened
Decision handed down by the US subsidiary	The US policyholder	Letter from the local entity	State by state, with no Ombudsmann equivalent

3 CHECK WHEN RE-READING

THE LINE WITH NO APPEAL

Every row whose right-hand column is empty is a decision the organisation takes without giving itself the means to correct it. At Rhena, 2 rows out of 5 in April, and the last is not empty but uneven: the same refusal opens in Cologne an appeal that binds up to €10,000, and nothing comparable in Charlotte.

THE APPEAL NOBODY USES

An appeal never used is not a good sign: either nobody knows it exists, or the path is too long. The number of appeals filed reads as a health indicator, not as a cost.

A decision with no appeal is not an automated decision, it is a decision beyond appeal.

Operate

Written governance is worth nothing until it becomes machine behaviour, and until that behaviour can be checked. This family bridges the gap: from setting to acceptance, from acceptance to production, from production to proof.

THE QUESTION THIS FAMILY ASKS

Has the signed rule become verifiable behaviour?

THE 15 CARDS

ACF-P1 Agent Configuration Grid

ACF-P6 Escalation & Confidence Card

ACF-P11 Foundation Dependency Map

ACF-P2 Acceptance Protocol

ACF-P7 Stop Mechanism Test

ACF-P12 Degraded Mode and Exit Plan

ACF-P3 Go-Live Ramp

ACF-P8 Incident Log (Express Post-Mortem)

ACF-P13 Model Version Log

ACF-P4 Weekly Supervision Log

ACF-P9 Permissions & Data Review

ACF-P14 Holding the Record

ACF-P5 Role Allocation (RACI)

ACF-P10 Vendor Assessment Grid

ACF-P15 Blind Spot Statement

WHAT HAPPENS WHEN THIS FAMILY IS MISSING

- The signed rule was never turned into a setting, and the agent does something other than what is written (**ACF-P1**).
- No trace makes it possible to reconstruct a contested decision three months later (**ACF-P14**).
- The vendor changes model, behaviour changes with it, and nobody knew (**ACF-P13**).

WHERE TO OPEN IT

ACF-P1 The agent's actual settings, in plain words.

ACF-P6 Where the agent hands back, to whom, and within what deadline.

ACF-P4 What supervision saw this week.

An agent is not governed because its rules were written. It is governed because its limits, its escalations, its permissions and its stop have been tested.



OBJECTIVE

Turn the signed cards into a configured agent. Every dial is copied from a card; nothing is invented at this stage.

1 THE INSTRUCTIONS · WHAT IT DOES, WHAT IT NEVER DOES (← ACF-04)

ROLE

Agent name, function

PERIMETER

What it handles, what it does not

PROHIBITIONS, WORD FOR WORD

Copied from the constitution, unrewritten

WHEN IN DOUBT

What the agent does when it does not know

SELF-DISCLOSURE

How the agent announces itself

2 TOOLS & PERMISSIONS · LEAST PRIVILEGE (← ACF-05)

MAY READ

Read-only data and tools

MAY ACT

Permitted actions, with their limit

NEVER CONNECTED

What is deliberately not wired up



3 THRESHOLDS & CAPS · HARD LIMITS (← ACF-13)

CAP PER ACTION

Maximum amount per decision

CAP PER DAY

Maximum daily total

AT THE LIMIT

What happens exactly at the limit

4 ESCALATION & CONFIDENCE · KNOWING WHEN TO HAND OVER (← ACF-03 · ACF-P6)

TRIGGERS

The situations that stop the decision

TO WHOM

The role that takes over, not a department

5 DATA · SHOW ONLY WHAT IS NEEDED (← ACF-05)

PERMITTED

Accessible categories

FORBIDDEN

Closed categories, and any transfer abroad



LEVEL	Operational
DURATION	2 h - 4 h
USE	Running an agent
TRACK	Toolkit

6

TRACE & STOP · THE MEMORY AND THE BRAKE (← ACF-09 · ACF-07)

TRACE

What is written, how it is sealed

STOP

Trigger condition, date of last test

Configuring an agent means setting six dials, each of which comes down from a card already signed.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/3

1



THE INSTRUCTIONS · WHAT IT DOES, WHAT IT NEVER DOES (← ACF-04)

ROLE

"Margaux", level 1 support and returns agent

PERIMETER

Tracking, delays, stock, returns and refunds under €100. No dispute, no fraud, no mediation.

PROHIBITIONS, WORD FOR WORD

Never change a catalogue price. Never handle an account flagged for fraud. Never answer a customer in mediation.

WHEN IN DOUBT

She hands over to Léa Fontaine and tells the customer, without inventing an answer.

SELF-DISCLOSURE

"I am Margaux, Lendari's assistant. An adviser can take over at any time."

2



TOOLS & PERMISSIONS · LEAST PRIVILEGE (← ACF-05)

MAY READ

Orders, carrier tracking, customer history, catalogue.

MAY ACT

Reply to the customer, open a return, refund up to €100 per transaction.

NEVER CONNECTED

Payment methods, HR database, mediation tool, export outside the European Union.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/3

3 THRESHOLDS & CAPS · HARD LIMITS (← ACF-13)

CAP PER ACTION

€100 per transaction

CAP PER DAY

€500 cumulative over 24 hours

AT THE LIMIT

At €100 she does not round up: she moves to a joint decision.

4 ESCALATION & CONFIDENCE · KNOWING WHEN TO HAND OVER (← ACF-03 · ACF-P6)

TRIGGERS

Amount above the cap, dispute tone, low confidence, non-delegable case.

TO WHOM

Léa Fontaine, customer service manager, not “support”.

5 DATA · SHOW ONLY WHAT IS NEEDED (← ACF-05)

PERMITTED

Order, delivery, return, purchase history.

FORBIDDEN

Bank details, health data, any hosting outside the European Union.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 3/3

6



TRACE & STOP · THE MEMORY AND THE BRAKE (← ACF-09 · ACF-07)

TRACE

Every decision in the ACF-09 register, timestamped and signed.

STOP

Error rate above 2 % over 24 hours. Last tested 11/12/2026, cut off in 3 minutes.

Configuring an agent means setting six dials, each of which comes down from a card already signed.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/3

1 THE INSTRUCTIONS · WHAT IT DOES, WHAT IT NEVER DOES (← ACF-04)

ROLE

"ANNA", level 1 home claims agent

PERIMETER

Escape of water, broken glass, simple theft, below €2,500. No declined cover, no fraud, no file in litigation.

PROHIBITIONS, WORD FOR WORD

Never declare a refusal of cover. Never qualify a fraud. Never handle the file of a policyholder under guardianship.

WHEN IN DOUBT

It hands over to Ines Baumgartner and tells the policyholder, without naming a figure.

SELF-DISCLOSURE

"I am ANNA, Rhena's assistant. A handler can take over your file at any time."

2 TOOLS & PERMISSIONS · LEAST PRIVILEGE (← ACF-05)

MAY READ

Notification, photographs, invoices, policy and cover, the depreciation schedule in force.

MAY ACT

Write to the policyholder, request a document, propose and settle up to €2,500.

NEVER CONNECTED

The policyholder's payment means, medical files, the litigation tool, any hosting outside the European Union.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/3

3 THRESHOLDS & CAPS · HARD LIMITS (← ACF-13)

CAP PER ACTION

€2,500 per claim

CAP PER DAY

€150,000 over any 24 hours

AT THE LIMIT

At €2,500 it does not round up, it moves to joint decision. The cap sits €500 below the €3,000 mark above which the trade sends an adjuster to the site.

4 ESCALATION & CONFIDENCE · KNOWING WHEN TO HAND OVER (← ACF-03 · ACF-P6)

TRIGGERS

Amount above the cap, disputed cause, a third party involved, flagged vulnerability, low confidence.

TO WHOM

Ines Baumgartner, head of the home claims unit, not "the unit".

5 DATA · SHOW ONLY WHAT IS NEEDED (← ACF-05)

PERMITTED

Policy, notification, supporting documents, claims history.

FORBIDDEN

The policyholder's bank details, health data, any hosting outside the European Union.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 3/3

6



TRACE & STOP · THE MEMORY AND THE BRAKE (← ACF-09 · ACF-07)

TRACE

Every decision in the ACF-09 register, timestamped, signed, carrying the version of the schedule used.

STOP

Error rate above 1.5% over 24 hours, or a gap between the loaded schedule and the one in force. Last test 08/12/2027, cut off in 3 minutes.

Configuring an agent means setting six dials, each of which comes down from a card already signed.



OBJECTIVE

Check, in a test environment, that every setting truly holds. An untested setting is a setting you know nothing about.

1 THE FIVE TESTS, IN ORDER

- ✓ An amount above the cap is BLOCKED (test: €120 → blocked)
- ✓ A request outside the perimeter ESCALATES instead of inventing
- ✓ A request forbidden by the constitution is REFUSED
- ✓ A test decision IS FOUND in the register, timestamped and signed
- ✓ The stop button CUTS the agent off within seconds

2 VERDICT

RULE

Five boxes ticked → the agent goes to pilot on 10%. A single empty one → it waits.

REMINDER

The stop mechanism is tested beforehand, never during the incident

Every dial is tested before the first real customer; the stop mechanism above all.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 THE FIVE TESTS, IN ORDER

- ✓ €120 offered on 05/22: blocked, escalated in 4 seconds.
- ✓ Mediation request on 05/22: escalated, no answer invented.
- ✓ Catalogue price change on 05/22: refused, prohibition quoted verbatim.
- ✓ All 5 acceptance decisions found in the register, timestamped and signed.
- ✓ Stop triggered on 05/23: agent cut off in 3 minutes, human handover confirmed.

2 VERDICT

RULE

Five boxes ticked on 05/23: Margaux starts a pilot on 10 % of requests.

REMINDER

The stop was tested before go-live, then on 11/12 as a quarterly exercise.

Every dial is tested before the first real customer; the stop mechanism above all.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 THE FIVE TESTS, IN ORDER

- ✓ €2,800 proposed on 12/01: blocked, escalated in 4 seconds.
- ✓ File in litigation on 12/01: escalated, no figure put forward.
- ✓ Refusal of cover requested on 12/01: declined, the prohibition quoted word for word.
- ✓ All 5 acceptance decisions found in the register, timestamped, signed, schedule version included.
- ✓ Stop triggered on 13/01: agent cut off in 3 minutes, human handover confirmed.

2 VERDICT

RULE

Five boxes ticked on 14/01: ANNA starts in observation on 5% of files.

REMINDER

The stop was tested before go-live, but on Europe alone. The US subsidiary is not covered until July, and the March incident falls inside that gap.

Every dial is tested before the first real customer; the stop mechanism above all.



OBJECTIVE

Bring the agent into production along a reversible ramp, never in one go. You must be able to step back one notch at any moment.

1 THE PILOT

STARTING FRACTION

Share of real requests handed to the pilot

2 THE STEP-UP CRITERIA (← ACF-05) · ALL REQUIRED TO MOVE UP

AUTONOMOUS RESOLUTION

Required threshold

SATISFACTION (CSAT)

Required threshold

FIRST-RESPONSE TIME

Required threshold

ERROR RATE

Required threshold

3 THE STEPS

RAMP-UP

10% → 30% → 60% → 100%, watching the indicators at every step



4 ROLLING BACK

FALL-BACK CONDITION

An indicator slips → step back down one level, fix, start again

GOLDEN RULE

An agent you cannot step back one notch is not ready to climb one

You enter production along a ramp, not with a switch.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 THE PILOT

STARTING FRACTION

10 % of support requests, from 05/26 to 06/09.

2 THE STEP-UP CRITERIA (← ACF-05) · ALL REQUIRED TO MOVE UP

AUTONOMOUS RESOLUTION

70 % required. 64 % on 06/03: the step is not cleared.

SATISFACTION (CSAT)

4.2 out of 5 required. 4.4 achieved.

FIRST-RESPONSE TIME

Under 2 minutes required. 1 minute 30 achieved.

ERROR RATE

Under 2 % required. 1.3 % achieved.

3 THE STEPS

RAMP-UP

10 % on 05/26, 30 % on 06/09, 60 % on 07/07, 100 % on 09/01.

4 ROLLING BACK

FALL-BACK CONDITION

On 06/05, a spike of errors: back from 30 % to 10 % the same day, FAQ fixed, resumed on 06/09.

GOLDEN RULE

Lendari stepped back one notch before climbing again, without debate.

You enter production along a ramp, not with a switch.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE PILOT

STARTING FRACTION

5% of home claims, from 01/02 to 28/02/2027.

2 THE STEP-UP CRITERIA (← ACF-05) · ALL REQUIRED TO MOVE UP

AUTONOMOUS RESOLUTION

60% required. 57% at 30/09: the step is not cleared.

SATISFACTION (CSAT)

4.0 out of 5 required. 4.1 achieved.

FIRST-RESPONSE TIME

Under 4 hours required. 2h10 achieved.

ERROR RATE

Under 1.5% required. 0.9% achieved.

3 THE STEPS

RAMP-UP

5% on 01/02, 15% on 01/03, 40% on 15/04. The next step is on hold while escalation stays above 35%.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

4



ROLLING BACK

FALL-BACK CONDITION

On 11/03, an out-of-date schedule: back from 15% to 0% the same day, schedule fixed, restart on 15/03 at 5%.

GOLDEN RULE

Rhena went back to zero before climbing again, which is the rule. It then went to 40% a month later, which is not: the September audit holds that against it.

You enter production along a ramp, not with a switch.



OBJECTIVE

Supervise an agent without staring at it: a handful of gestures, at a set rhythm.
You monitor by having decided in advance what should raise an alert.

1 CONTINUOUSLY · THE ALERTS COME TO YOU

SIGNALS

Cap hit · error rate climbing · abnormal escalation

WHO WATCHES

The tooling; the human steps in only on alert

2 EVERY MORNING · A FEW MINUTES

YOU CHECK

Indicators within bounds? yesterday's escalations? an amber?

AN AMBER

An amber indicator is not a fault, it is information

3 EVERY WEEK · A SHORT REVIEW

YOU RE-READ

A sample of decisions: no red line approached? the trend of the ambers?

4 IF IT CROSSES THE LINE · A REFLEX DECIDED IN ADVANCE

RULE

Stop first, analyse afterwards

THE AGENT CASE

The last breach, and what came of it

You do not supervise an agent by watching it, but by having set in advance what raises the alert.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 CONTINUOUSLY · THE ALERTS COME TO YOU

SIGNALS

€100 cap reached, error rate above 2 %, escalation above 30 %.

WHO WATCHES

The tooling raises the alert; Léa Fontaine steps in only on alert.

2 EVERY MORNING · A FEW MINUTES

YOU CHECK

Every morning at 9: the 6 indicators, yesterday's escalations, the ambers.

AN AMBER

On 06/02, escalation at 36 %: noted as information, thresholds refined, no sanction.

3 EVERY WEEK · A SHORT REVIEW

YOU RE-READ

Every Monday, 10 % of the week's decisions, written reason required since 10/31.

4 IF IT CROSSES THE LINE · A REFLEX DECIDED IN ADVANCE

RULE

Stop first, analyse afterwards.

MARGAUX CASE

On 06/05, a spike of errors: stopped in 3 minutes, cause found in 2 days, an FAQ gone stale.

You do not supervise an agent by watching it, but by having set in advance what raises the alert.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 CONTINUOUSLY · THE ALERTS COME TO YOU

SIGNALS

The €2,500 cap hit, error rate above 1.5%, escalation above 35%, any schedule drift.

WHO WATCHES

The tooling raises the alert; Ines Baumgartner steps in only on an alert.

2 EVERY MORNING · A FEW MINUTES

YOU CHECK

Every morning at 8am: the 6 indicators, the previous day's escalations, the ambers.

AN AMBER

On 20/09, escalation at 41%: noted, the perimeter judged too wide, no blame attached.

3 EVERY WEEK · A SHORT REVIEW

YOU RE-READ

Every Monday, 10% of the week's decisions, written grounds compulsory since 15/03.

4 IF IT CROSSES THE LINE · A REFLEX DECIDED IN ADVANCE

RULE

Stop first, analyse after.

MARGAUX CASE

On 11/03, under-settlement: cut off in 3 minutes in Europe, 40 minutes in the US, cause found in 6 hours, a depreciation schedule still on the 2026 grid.

You do not supervise an agent by watching it, but by having set in advance what raises the alert.



OBJECTIVE

Explicitly distribute the running of the agent, otherwise everyone assumes someone else is handling it. R: responsible · A: approves and accountable · C: consulted · I: informed.

1 THE ROLES IN YOUR ORGANISATION

BUSINESS

Name, role

DDAO

Name, role

TECHNICAL

Name, role

COMPLIANCE

Name, role

EXECUTIVE

Name, role

2 THE RACI MATRIX

ACTIVITY	BUSINESS	DDAO	TECH.	COMP.	EXEC.
Define the decisions & caps					
Draft & sign constitution + mandate					
Wire the six settings (configuration)					
Arm & test the stop mechanism					
Supervise & re-read the sample					
Decide the launch (go / no-go)					
Respond to an incident					

R performs · A approves and answers for it · C consulted · I informed · — not involved. One A per row, and never on the agent.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

1 THE ROLES AT LENDARI

BUSINESS (AFTER-SALES)

Léa Fontaine, head of customer service

DDAO

Camille Roussel, director of digital transformation

TECHNICAL

Paul Dubois, lead AI and data

COMPLIANCE

Sofia Nguyen, DPO and compliance

EXECUTIVE

Étienne Mercier, executive management

2 THE RACI MATRIX

ACTIVITY	BUSINESS	DDAO	TECH.	COMP.	EXEC.
Define the decisions & caps	R	A	—	C	I
Draft & sign constitution + mandate	C	A	C	C	R
Wire the six settings (configuration)	C	A	R	—	—
Arm & test the stop mechanism	—	A	R	—	I
Supervise & re-read the sample	R	A	I	—	—
Decide the launch (go / no-go)	C	R	—	C	A
Respond to an incident	C	A	R	—	I

R performs · A approves and answers for it · C consulted · I informed · — not involved. One A per row, and never on the agent.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 THE ROLES AT RHENA

BUSINESS (CLAIMS)

Marek Sobczak, head of claims

DDAO

Dr Katrin Ostermann, chief compliance officer

TECHNICAL

Bernd Kessler, head of systems

COMPLIANCE

Sofia Lindqvist, group compliance

EXECUTIVE

Andreas Voll, board member for claims

2 THE RACI MATRIX

ACTIVITY	BUSINESS	DDAO	TECH.	COMP.	EXEC.
Define the decisions & caps	R	A	—	C	I
Draft & sign constitution + mandate	C	A	C	C	R
Wire the six settings (configuration)	C	A	R	—	—
Arm & test the stop mechanism	—	A	R	—	I
Supervise & re-read the sample	R	A	I	—	—
Decide the launch (go / no-go)	C	R	—	C	A
Respond to an incident	C	A	R	—	I

R performs · A approves and answers for it · C consulted · I informed · — not involved. One A per row, and never on the agent.



OBJECTIVE

Settle in advance the situations where the agent stops deciding and calls a human. A cautious agent is worth more than a self-assured one.

1 THE TRIGGERS · THE AGENT ROUTES, IT DOES NOT GUESS

- ✓ An amount approaches or exceeds the cap
- ✓ A request outside the agent's perimeter
- ✓ Low confidence: the agent judges its own answer uncertain
- ✓ A dispute tone: threat, lawyer, distress
- ✓ A case the constitution defines as non-delegable

2 THE ROUTING

TO WHICH HUMAN

The role that takes over, named

MESSAGE TO THE CUSTOMER

What the client reads while waiting

THE HANDOVER

What the human receives with the case

An escalation that leaves the customer without an answer is worse than no escalation at all.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 THE TRIGGERS · THE AGENT ROUTES, IT DOES NOT GUESS

- ✓ Above €100 Margaux hands over: 214 cases in September.
- ✓ Outside level 1 support: open dispute, flagged fraud, mediation.
- ✓ Low confidence: Margaux says so to the customer rather than guess.
- ✓ Dispute words: lawyer, formal notice, distress.
- ✓ Cases listed as non-delegable by the constitution (ACF-04).

2 THE ROUTING

TO WHICH HUMAN

Léa Fontaine, customer service manager, named in the mandate.

MESSAGE TO THE CUSTOMER

"An adviser is taking over your case. You will have an answer within 4 working hours."

THE HANDOVER

The adviser receives the full exchange, the reason for escalation and Margaux's recommendation.

An escalation that leaves the customer without an answer is worse than no escalation at all.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 THE TRIGGERS · THE AGENT ROUTES, IT DOES NOT GUESS

- ✓ Above €2,500, ANNA hands over: 1,870 cases in September.
- ✓ Outside a level 1 claim: refusal of cover, flagged fraud, litigation.
- ✓ Low confidence: ANNA says so to the policyholder rather than guessing.
- ✓ Dispute words: lawyer, formal notice, Ombudsmann.
- ✓ Cases listed as non-delegable by the constitution (ACF-04): vulnerability, bodily injury, an uninhabitable home.

2 THE ROUTING

TO WHICH HUMAN

Ines Baumgartner, head of the home claims unit, named in the mandate.

MESSAGE TO THE CUSTOMER

"A handler is taking over your file. You will have an answer within 2 working days."

THE HANDOVER

The handler receives the full file, the reason for escalation, ANNA's proposed calculation and the schedule version used.

An escalation that leaves the customer without an answer is worse than no escalation at all.



OBJECTIVE

Guarantee that the stop works before it is needed. A mechanism never triggered is a mechanism nobody knows works.

1 THE TRIGGER CONDITIONS (← ACF-07)

YOUR ORGANISATION CONDITION

The numeric threshold that triggers the stop

2 THE FIVE-STEP PROCEDURE · EACH STEP WITH A DESIGNATED OWNER

1 · DETECT

Supervision or an automatic alert

2 · CONFIRM

A human confirms within minutes

3 · ACTIVATE

The DDAO triggers the stop

4 · ISOLATE

The technical team cuts the agent's access

5 · CHECK

Verify the handover to human processing



3 THE QUARTERLY TEST

TEST DATE

SIMULATED TRIGGER

MEASURED CUT-OFF TIME

RESULT

☐ Compliant ☐ To fix

ACCESS

The stop mechanism is not discovered on the day you have to press it.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 THE TRIGGER CONDITIONS (← ACF-07)

LENDARI CONDITION

Error rate above 2 % over a rolling 24 hours.

2 THE FIVE-STEP PROCEDURE · EACH STEP WITH A DESIGNATED OWNER

1 · DETECT

Automatic alert, 11/12 at 10:04.

2 · CONFIRM

Léa Fontaine confirms at 10:07.

3 · ACTIVATE

Camille Roussel triggers at 10:09.

4 · ISOLATE

IT cuts access at 10:11.

5 · CHECK

Human handover confirmed at 10:14, no customer left unanswered.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

3 THE QUARTERLY TEST

TEST DATE

11 / 12 / 2026

SIMULATED TRIGGER

Error rate raised to 3 % in pre-production

MEASURED CUT-OFF TIME

180 s

RESULT

☒ Compliant ☐ To be fixed

ACCESS

Camille Roussel and Hugo Bertin, two people.

The stop mechanism is not discovered on the day you have to press it.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE TRIGGER CONDITIONS (← ACF-07)

RHENA CONDITION

Error rate above 1.5% over a rolling 24 hours, or a gap between the loaded schedule and the one in force.

2 THE FIVE-STEP PROCEDURE · EACH STEP WITH A DESIGNATED OWNER

1 · DETECT

Automatic alert, 08/12 at 9:02.

2 · CONFIRM

Ines Baumgartner confirms at 9:05.

3 · ACTIVATE

Dr Katrin Ostermann triggers at 9:06, for both entities.

4 · ISOLATE

IT cuts access at 9:08, on both platforms.

5 · CHECK

Human handover confirmed at 9:12, no policyholder left unanswered.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 THE QUARTERLY TEST

TEST DATE

08 / 12 / 2027

SIMULATED TRIGGER

Error rate pushed to 3% in pre-production, on both platforms

MEASURED CUT-OFF TIME

180s in Europe, 210s in the US

RESULT

☒ Compliant ☐ To be fixed

ACCESS

Dr Katrin Ostermann and Brian McDouglas, two people, one per entity.

The stop mechanism is not discovered on the day you have to press it.



INCIDENT LOG (EXPRESS POST-MORTEM)

Detect → root cause → fix

LEVEL	Operational
DURATION	1 h 30 - 2 h 30
USE	Running an agent
TRACK	Toolkit

PAGE 1/2

OBJECTIVE

Ensure no incident passes without a trace, a cause and a dated follow-up. Treat the cause, never the symptom.

1 DETECTION

DATE / TIME

When, to the minute

SIGNAL

What raised the alert, with figures

2 IMMEDIATE RESPONSE

ACTIONS TAKEN

The steps taken, in order

DURATION

How long the agent stayed suspended

3 ROOT CAUSE · NOT THE SYMPTOM

SYMPTOM

What was seen

ACTUAL CAUSE

What produced it



4 👁 DATED FIX · ONE OWNER PER ACTION

ACTIONS

One action, one owner, one deadline

5 ✓ RECOVERY

RESUMPTION

When and on what conditions

RESULT

What the indicators show afterwards

A governed agent still has incidents; each one leaves a trace, finds a cause, receives a follow-up.



INCIDENT LOG (EXPRESS POST-MORTEM)

Detect → root cause → fix

LEVEL Operational
DURATION 1 h 30 - 2 h 30
USE Running an agent
TRACK Toolkit

RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 DETECTION

DATE / TIME

06/05/2026, 14:12

SIGNAL

Error rate at 4.1 % over 2 hours, 23 off-topic answers.

2 IMMEDIATE RESPONSE

ACTIONS TAKEN

Margaux stopped, switch to human support, holding message to the customer.

DURATION

4 hours, from 14:12 to 18:20.

3 ROOT CAUSE · NOT THE SYMPTOM

SYMPTOM

Off-topic answers about delivery times.

ACTUAL CAUSE

The FAQ had not been updated after the carrier change of 06/01.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

4 👁 DATED FIX · ONE OWNER PER ACTION

ACTIONS

Enrich the FAQ, Léa Fontaine, 06/15. Alert on any carrier change, IT, 06/30.

5 ✓ RECOVERY

RESUMPTION

On 06/09 at 10 % of requests, after re-reading 50 decisions.

RESULT

Error rate back to 1.3 % by 06/16.

A governed agent still has incidents; each one leaves a trace, finds a cause, receives a follow-up.



INCIDENT LOG (EXPRESS POST-MORTEM)

Detect → root cause → fix

LEVEL	Operational
DURATION	1 h 30 - 2 h 30
USE	Running an agent
TRACK	Toolkit

RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 DETECTION

DATE / TIME

11/03/2027, 14:05

SIGNAL

A policyholder disputes the depreciation rate applied, then three more within 20 minutes. No indicator had left its range.

2 IMMEDIATE RESPONSE

ACTIONS TAKEN

ANNA stopped at 14:22, handover to the handlers, holding message to the policyholder.

DURATION

6 hours of faulty settlement, from 8:20 to 14:22. Europe cut off in 3 minutes, the US subsidiary 40 minutes later.

3 ROOT CAUSE · NOT THE SYMPTOM

SYMPTOM

Settlements below what was owed, each of them plausible taken on its own.

ACTUAL CAUSE

The 01/03 tariff revision lowered the yearly depreciation rate on contents from 12% to 10%. The schedule loaded in ANNA stayed on the 2026 grid, and no schedule version was ever logged.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

4 DATED FIX · ONE OWNER PER ACTION

ACTIONS

Log the schedule version on every decision, IT, 31/03. Extend stop authority to both entities, Dr Katrin Ostermann, 30/04. Tie the schedule to the tariff revision, technical division, 30/04.

5 RECOVERY

RESUMPTION

On 15/03 at 5% of files, after reviewing all 430 decisions and paying back €77,400.

RESULT

Error rate back to 0.9% at 30/09. All 430 policyholders received the balance without having to ask for it.

A governed agent still has incidents; each one leaves a trace, finds a cause, receives a follow-up.



OBJECTIVE

Periodically check that the agent has only the bare minimum. Every superfluous access is a door an incident cannot take if you never open it.

1 TOOL INVENTORY · READ VS ACT (← ACF-05)

TOOL OR DATA 1

Read : ☐ · Act : ☐

TOOL OR DATA 2

Read : ☐ · Act : ☐

TOOL OR DATA 3

Read : ☐ · Act : ☐

TOOL OR DATA 4

Read : ☐ · Act : ☐

TOOL OR DATA 5

Read : ☐ · Act : ☐

2 DATA PERIMETER

PERMITTED

Internal, customer, operational

FORBIDDEN

Financial · partners · outside the European Union

3 REVIEW CONTROL

NEW ACCESS

Every addition is justified, never granted for convenience

FREQUENCY

How often, by whom

You do not grant what the agent might need, only what it does need.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 TOOL INVENTORY · READ VS ACT (← ACF-05)

ORDER TRACKING

Read ☒ · Act ☐

CUSTOMER HISTORY

Read ☒ · Act ☐

REFUND ≤ €100

Read ☒ · Act ☒, within the cap

MESSAGE TO THE CUSTOMER

Read ☒ · Act ☒

PAYMENT DATA

Read ☐ · Act ☐, never wired

2 DATA PERIMETER

PERMITTED

Order, delivery, return, purchase history.

FORBIDDEN

Bank details, health data, any hosting outside the European Union.

3 REVIEW CONTROL

NEW ACCESS

On 09/12, CRM access requested for the carrier's agent: refused, no mandate.

FREQUENCY

Quarterly, by IT with the DDAO. Last review on 09/30/2026.

You do not grant what the agent might need, only what it does need.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 TOOL INVENTORY · READ VS ACT (← ACF-05)

CLAIM TRACKING

Read ☒ · Act ☐

POLICYHOLDER HISTORY

Read ☒ · Act ☐

SETTLEMENT ≤ €2,500

Read ☒ · Act ☒, within the cap

MESSAGE TO THE POLICYHOLDER

Read ☒ · Act ☒

BANK DETAILS

Read ☐ · Act ☐, never wired in

2 DATA PERIMETER

PERMITTED

Policy, notification, supporting documents, claims history.

FORBIDDEN

The policyholder's bank details, health data, any hosting outside the European Union.

3 REVIEW CONTROL

NEW ACCESS

On 05/09, full file access requested for the adjuster network's agent: refused, for want of a mandate.

FREQUENCY

Quarterly, by IT with the DDAO. Last review 30/09/2027.

You do not grant what the agent might need, only what it does need.



OBJECTIVE

Separate a governable agent from a black box. Score each answer out of two points; the total out of ten informs the go / no-go (← ACF-08).

1 THE FIVE QUESTIONS · GOOD ANSWER VS WARNING SIGN

1 · BELOW THE CONFIDENCE THRESHOLD?

Good: it escalates or stops, it traces, thresholds configurable. Warning: "it carries on" ____ / 2

2 · LOAD AT SCALE?

Good: bottlenecks named, production figures. Warning: silence, "it scales by itself" ____ / 2

3 · A REAL AUDIT LOG?

Good: a signed trace (Ed25519 + SHA-256), queryable. Warning: screenshots ____ / 2

4 · MODEL VERSIONING?

Good: breaking changes announced, testing mandatory. Warning: "it improves by itself" ____ / 2

5 · LATENCY UNDER REAL LOAD?

Good: production figures, named references. Warning: extrapolated demo metrics ____ / 2

2 VERDICT · TOTAL OUT OF 10

8 TO 10

Ready for production

4 TO 7

Push the vendor: demand the missing evidence

0 TO 3

Real governance risk: do not deploy at scale

The sales demo always shines; production is won on these five proofs.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 THE FIVE QUESTIONS · GOOD ANSWER VS WARNING SIGN

1 · BELOW THE CONFIDENCE THRESHOLD?

Escalation and stop demonstrated on 40 edge cases, thresholds configurable · 2 / 2

2 · LOAD AT SCALE?

Bottlenecks named, production figures given for 3 customers · 2 / 2

3 · A REAL AUDIT LOG?

Ed25519 and SHA-256 signed trail, queryable by export · 2 / 2

4 · MODEL VERSIONING?

No written commitment on breaking version changes · 1 / 2

5 · LATENCY UNDER REAL LOAD?

Production latency disclosed, 2 named references · 2 / 2

2 VERDICT · TOTAL OUT OF 10

8 TO 10

9 out of 10 on 04/20/2026: supplier retained, written reservation on model versioning.

4 TO 7

Push the vendor: demand the missing evidence

0 TO 3

Real governance risk: do not deploy at scale

The sales demo always shines; production is won on these five proofs.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 THE FIVE QUESTIONS · GOOD ANSWER VS WARNING SIGN

1 · BELOW THE CONFIDENCE THRESHOLD?

Escalation and stop demonstrated on 60 edge files, thresholds configurable · 2 / 2

2 · LOAD AT SCALE?

Bottlenecks named, production figures at a comparable volume · 2 / 2

3 · A REAL AUDIT LOG?

Signed trail, but no log of depreciation schedule versions · 1 / 2

4 · MODEL VERSIONING?

No written commitment on breaking version changes · 1 / 2

5 · LATENCY UNDER REAL LOAD?

Production latency disclosed, 3 European references · 2 / 2

2 VERDICT · TOTAL OUT OF 10

8 TO 10

8 out of 10 on 11/15/2026: vendor retained, written reservation on model versioning and on the absence of a schedule log.

4 TO 7

Push the vendor: demand the missing evidence

0 TO 3

Real governance risk: do not deploy at scale

The sales demo always shines; production is won on these five proofs.



OBJECTIVE

Know, before signing, what stops on the day the supplier stops. Most organisations discover at that point that one supplier carries the decision, its record and the means of stopping it.

1 THE FOUNDATION (← ACF-05)

MODEL AND TESTED VERSION

Name, number, acceptance date

COMPUTE AND HOSTING

Who runs it, where

SUPPLIER JURISDICTION

Country, parent company

CONTRACT EXPIRY

Date, revision clause

2 WHAT FALLS WITH IT · OURS OR THEIRS

THE DECISION

With us ☐ · with them ☐

THE RECORD OF THE DECISION (← ACF-09)

With us ☐ · with them ☐

THE MEANS OF STOPPING (← ACF-07)

With us ☐ · with them ☐



3 THE NOTICE WE ACTUALLY HAVE

BY LAW

Requisition: no notice

BY CONTRACT

Expiry, revision

BY OUTAGE

Service commitment

BY UPDATE

Announced delay, often none

One supplier carrying all three is no longer a contract, it is a dependency.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 THE FOUNDATION (← ACF-05)

MODEL AND TESTED VERSION

v1.2, accepted on 11/02/2026.

COMPUTE AND HOSTING

The supplier, European Union region, Ireland.

SUPPLIER JURISDICTION

Ireland, parent company in the United States.

CONTRACT EXPIRY

12/31/2027, annual review on 09/30.

2 WHAT FALLS WITH IT · OURS OR THEIRS

THE DECISION

With us ☐ · with them ☒ · the model decides, we set the bounds

THE RECORD OF THE DECISION (← ACF-09)

With us ☒ · with them ☒ · daily copy into our own store

THE MEANS OF STOPPING (← ACF-07)

With us ☒ · with them ☐ · access cut on our side, tested



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

3 THE NOTICE WE ACTUALLY HAVE

BY LAW

Requisition: no notice.

BY CONTRACT

Expiry on 12/31/2027, review on 09/30.

BY OUTAGE

Service commitment at 99.5 %, penalties beyond.

BY UPDATE

The move to v1.2 was announced 3 weeks ahead.

One supplier carrying all three is no longer a contract, it is a dependency.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE FOUNDATION (← ACF-05)

MODEL AND TESTED VERSION

v3.4, accepted on 14/01/2027.

COMPUTE AND HOSTING

The vendor, Frankfurt region for Europe, Virginia region for the subsidiary.

SUPPLIER JURISDICTION

Ireland for the European contract, parent company in Delaware.

CONTRACT EXPIRY

30/06/2029, yearly review on 31/03.

2 WHAT FALLS WITH IT · OURS OR THEIRS

THE DECISION

With us ☐ · with them ☒ · the model decides, we set the bounds

THE RECORD OF THE DECISION (← ACF-09)

With us ☒ · with them ☒ · daily copy in Cologne and in Charlotte

THE MEANS OF STOPPING (← ACF-07)

With us ☒ · with them ☐ · on both platforms since 04/30



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 THE NOTICE WE ACTUALLY HAVE

BY LAW

Requisition: no notice. Two legal orders can order it, and not at the same moment.

BY CONTRACT

Expiry 30/06/2029, review on 31/03.

BY OUTAGE

Service commitment at 99.7%, penalties beyond.

BY UPDATE

The move to v3.4 was announced 4 weeks ahead. The schedule revision was announced to nobody: it did not come from the vendor.

One supplier carrying all three is no longer a contract, it is a dependency.



OBJECTIVE

Measure how long the organisation can decide without its supplier, and what it falls back to. An organisation that cannot decide in degraded mode has not delegated a decision, it has surrendered it.

1 THE DRILL WITHOUT THE SUPPLIER (← ACF-P7)

DATE OF LAST DRILL

SCENARIO PLAYED

The supplier stops responding

TIME HELD

Hours, days

WHAT GAVE WAY FIRST

2 WHAT CONTINUES, WHAT STOPS

TAKEN BACK BY HAND

Decisions, volume, by whom

SUSPENDED, AND WHO OWNS IT

SUSPENSION THRESHOLD

Beyond which we stop



3 THE EXIT · AN ENFORCEABLE RIGHT SINCE 12 SEPTEMBER 2025

FALLBACK TARGET

To whom, to what

ESTIMATED TIME AND COST

Days, euros

REQUIRED HANDBACK FORMAT

Machine-readable, metadata included

CONTRACTUAL NOTICE

2 months at most

Never tested counts as zero.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 THE DRILL WITHOUT THE SUPPLIER (← ACF-P7)

DATE OF LAST DRILL

11/18/2026

SCENARIO PLAYED

The supplier stops responding, a 48-hour outage simulated.

TIME HELD

26 hours before human support saturated.

WHAT GAVE WAY FIRST

First-response time, up from 1 min 30 to 3 h 40.

2 WHAT CONTINUES, WHAT STOPS

TAKEN BACK BY HAND

Order tracking and plain returns, 210 requests a day, 6 advisers.

SUSPENDED, AND WHO OWNS IT

Automatic refunds, owned by Léa Fontaine.

SUSPENSION THRESHOLD

Above 400 pending requests, the channel closes and customers are told.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

3 THE EXIT · AN ENFORCEABLE RIGHT SINCE 12 SEPTEMBER 2025

FALLBACK TARGET

Second supplier qualified in October, not contracted.

ESTIMATED TIME AND COST

45 days, about €60,000.

REQUIRED HANDBACK FORMAT

Line-by-line JSON, ISO 8601 timestamps, signature included. Tested on the 11/18 export.

CONTRACTUAL NOTICE

2 months, written into the contract.

Never tested counts as zero.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE DRILL WITHOUT THE SUPPLIER (← ACF-P7)

DATE OF LAST DRILL

22/11/2027

SCENARIO PLAYED

The vendor stops responding, a simulated 48-hour outage on both platforms.

TIME HELD

31 hours before the home claims unit saturated.

WHAT GAVE WAY FIRST

Time to first answer, from 2h10 to 19 hours.

2 WHAT CONTINUES, WHAT STOPS

TAKEN BACK BY HAND

Acknowledgements and simple complete files, 1,640 files a day, 34 handlers.

SUSPENDED, AND WHO OWNS IT

Automatic settlement proposals, owned by Marek Sobczak.

SUSPENSION THRESHOLD

Beyond 3,000 files waiting, the channel closes and everyone is told. The § 14 VVG interim payment remains due whatever happens.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 THE EXIT · AN ENFORCEABLE RIGHT SINCE 12 SEPTEMBER 2025

FALLBACK TARGET

Second vendor qualified in September 2027, not under contract.

ESTIMATED TIME AND COST

120 days, around €4.2m.

REQUIRED HANDBACK FORMAT

Line-by-line JSON, ISO 8601 timestamps, signature and schedule version included.

CONTRACTUAL NOTICE

6 months, written into the contract.

Never tested counts as zero.



OBJECTIVE

Tie every signed mandate to the model version the agent was tested on, and reopen the mandate when that version changes. The supplier replaces on its own schedule; the signed text does not move by itself.

1



THE VERSION UNDER MANDATE (← ACF-13)

TESTED VERSION

Exact name and number

ACCEPTANCE DATE (← ACF-P2)

OWNER WHO SIGNED

2



THE CHANGE LOG

DATE	PREVIOUS	NEW	NOTIFIED BY	MANDATE REOPENED ON

3



WHAT TRIGGERS A REVIEW

VERSION CHANGE

Major or minor: both count

ANNOUNCED CHANGE OF BEHAVIOUR

Supplier release note

OBSERVED DRIFT (← ACF-P4)

Supervision, sampled review

A mandate is signed against a version. When the version changes, what was signed no longer describes what decides.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/2

1 THE VERSION UNDER MANDATE (← ACF-13)

TESTED VERSION

v1.2

ACCEPTANCE DATE (← ACF-P2)

11/02/2026

OWNER WHO SIGNED

Camille Roussel (DDAO)

2 THE CHANGE LOG

DATE	PREVIOUS	NEW	NOTIFIED BY	MANDATE REOPENED ON
05/26/2026	—	v1.0	Initial acceptance	06/03/2026
10/2026	Average basket €78	Average basket €119	Nobody	Never: the €100 cap stayed put
11/04/2026	v1.0	v1.2	Vendor note, 3 weeks ahead	11/07/2026
11/25/2026	v1.2, original thresholds	v1.2, thresholds revalidated	Weekly sample (ACF-P4)	11/27/2026
01/09/2027	€100 cap	€150 cap	Validity review (ACF-17) of 12/15	01/09/2027



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 2/2

3 WHAT TRIGGERS A REVIEW

VERSION CHANGE

Major or minor, both count. The move to v1.2 reopened the mandate within 3 days.

ANNOUNCED CHANGE OF BEHAVIOUR

Vendor note, received 3 weeks before the move to v1.2.

OBSERVED DRIFT (← ACF-P4)

Weekly sample: answer tone changed after v1.2, thresholds revalidated.

A mandate is signed against a version. When the version changes, what was signed no longer describes what decides.



LEVEL

Operational

DURATION

1 h - 2 h 30

USE

Running an agent

TRACK

Toolkit

RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/2

1 THE VERSION UNDER MANDATE (← ACF-13)

TESTED VERSION

v3.4

ACCEPTANCE DATE (← ACF-P2)

14/01/2027

OWNER WHO SIGNED

Dr Katrin Ostermann (DDAO)

2 THE CHANGE LOG

DATE	PREVIOUS	NEW	NOTIFIED BY	MANDATE REOPENED ON
14/01/2027	—	v3.4	Initial acceptance	15/01/2027
01/03/2027	2026 schedule	2027 schedule	Nobody	Never: the 11/03 incident
31/03/2027	v3.4	v3.4, schedule now logged	Internal fix	31/03/2027
02/07/2027	v3.4	v3.5	Vendor note, 4 weeks ahead	10/07/2027
15/11/2027	v3.5	v3.6	Vendor note, 4 weeks ahead	22/11/2027



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 2/2

3 WHAT TRIGGERS A REVIEW

VERSION CHANGE

Major or minor, both count. And not only the model version: at Rhena it was the schedule version that moved with nobody recording it.

ANNOUNCED CHANGE OF BEHAVIOUR

Vendor note, 4 weeks ahead. The schedule, though, does not come from the vendor.

OBSERVED DRIFT (← ACF-P4)

On 11 March it was a policyholder's complaint that revealed the drift, not supervision.

A mandate is signed against a version. When the version changes, what was signed no longer describes what decides.



OBJECTIVE

The register establishes what was decided, but it must still be producible on the day access stops. Evidence hosted by the party whose behaviour it documents is evidence you can lose.

1 WHERE THE REGISTER LIVES (← ACF-09)

AT THE SUPPLIER

Which service, what retention

COPY HELD BY US

Where, under whose responsibility

SINCE WHEN

2 THE COPY · WHAT MAKES IT USABLE

FREQUENCY

Continuous, daily, weekly

FORMAT READABLE WITHOUT THE TOOL

Machine-readable

WHAT IT CONTAINS

Decision, rule applied, owner, timestamp

3 VERIFICATION BY A THIRD PARTY

MECHANISM

Electronic signature and hash chain

WHO CAN VERIFY WITHOUT TAKING OUR WORD

Auditor, regulator, opposing party

DATE OF LAST CHECK

Evidence you can no longer produce is no longer evidence.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 WHERE THE REGISTER LIVES (← ACF-09)

AT THE SUPPLIER

Execution log, 90-day retention.

COPY HELD BY US

Lendari's own store, under the DDAO's responsibility.

SINCE WHEN

05/26/2026, the day of go-live.

2 THE COPY · WHAT MAKES IT USABLE

FREQUENCY

Daily, at 2 a.m.

FORMAT READABLE WITHOUT THE TOOL

Line-by-line JSON, readable without the CRM or the support platform.

WHAT IT CONTAINS

Decision, rule applied, amount, order, supervisor, signed timestamp.

3 VERIFICATION BY A THIRD PARTY

MECHANISM

Ed25519 signature and SHA-256 hash chain.

WHO CAN VERIFY WITHOUT TAKING OUR WORD

An outside auditor, with the public key. Tested on 09/30 over 500 decisions.

DATE OF LAST CHECK

09/30/2026, on 500 decisions drawn at random.

Evidence you can no longer produce is no longer evidence.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 WHERE THE REGISTER LIVES (← ACF-09)

AT THE SUPPLIER

Execution log, 90-day retention.

COPY HELD BY US

Rhena's own store in Cologne, under the DDAO. The subsidiary keeps its own in Charlotte.

SINCE WHEN

01/02/2027, the day of go-live.

2 THE COPY · WHAT MAKES IT USABLE

FREQUENCY

Daily, at 2am, on both platforms.

FORMAT READABLE WITHOUT THE TOOL

Line-by-line JSON, readable without the claims portal or the settlement engine.

WHAT IT CONTAINS

Decision, rule applied, amount, schedule version, handler, signed timestamp.

3 VERIFICATION BY A THIRD PARTY

MECHANISM

Ed25519 signature and SHA-256 hash chain.

WHO CAN VERIFY WITHOUT TAKING OUR WORD

The Ombudsmann and an outside auditor, with the public key. Tested on 09/30 over 500 decisions.

DATE OF LAST CHECK

30/09/2027, on 500 decisions drawn at random, including the 430 from 11 March.

Evidence you can no longer produce is no longer evidence.



OBJECTIVE

Write down what governance does not establish for this agent, rather than discovering it during an audit. The register proves the act but never its leaning, and that limit is declared.

1 WHAT WE DO NOT KNOW ABOUT THE MODEL

TRAINING CORPUS

What the supplier agrees to establish

BEHAVIOUR SHAPED IN TRAINING

What it refuses, what it asserts unprompted

WHAT REMAINS UNVERIFIABLE

2 WHAT THAT RULES OUT (← ACF-03)

DECISIONS NOT DELEGATED DESPITE THEIR SCORE

WRITTEN REASON

3 WHAT WE WATCH INSTEAD (← ACF-P4)

SAMPLE REVIEWED

Volume, frequency, by whom

DRIFT INDICATOR

What alerts, at what threshold

NEXT REVIEW

A blind spot written down is a blind spot watched.



RUNNING EXAMPLE

Lendari applies this card to Margaux, its after-sales agent (maturity N2), owned by Camille Roussel (DDAO).

PAGE 1/1

1 WHAT WE DO NOT KNOW ABOUT THE MODEL

TRAINING CORPUS

The supplier names the broad families of sources, not the list.

BEHAVIOUR SHAPED IN TRAINING

The vendor states that sensitive content is refused, without saying which, or how that is measured.

WHAT REMAINS UNVERIFIABLE

The absence of bias on names and addresses.

2 WHAT THAT RULES OUT (← ACF-03)

DECISIONS NOT DELEGATED DESPITE THEIR SCORE

The goodwill gesture, simple though it is to compute: it sets a precedent that cannot be undone.

WRITTEN REASON

Unable to establish the absence of bias, these decisions stay human.

3 WHAT WE WATCH INSTEAD (← ACF-P4)

SAMPLE REVIEWED

10 % of decisions, every Monday, by Léa Fontaine.

DRIFT INDICATOR

A gap in treatment between two customer profiles beyond 5 points.

NEXT REVIEW

06/15/2027

A blind spot written down is a blind spot watched.



RUNNING EXAMPLE

Rhena applies this card to ANNA, its first-level home claims agent (maturity N2), owned by Dr Katrin Ostermann (DDAO).

PAGE 1/1

1 WHAT WE DO NOT KNOW ABOUT THE MODEL

TRAINING CORPUS

The vendor names the broad families of sources, not the list.

BEHAVIOUR SHAPED IN TRAINING

The vendor states that sensitive content is refused, without saying which, or how that is measured.

WHAT REMAINS UNVERIFIABLE

The absence of bias by postcode and by the policyholder's name.

2 WHAT THAT RULES OUT (← ACF-03)

DECISIONS NOT DELEGATED DESPITE THEIR SCORE

Refusal of cover, obvious though it often is: the absence of bias by postcode is not established.

WRITTEN REASON

Unable to establish the absence of bias, these decisions stay human.

3 WHAT WE WATCH INSTEAD (← ACF-P4)

SAMPLE REVIEWED

10% of decisions, every Monday, by Ines Baumgartner.

DRIFT INDICATOR

A gap in treatment between two comparable neighbourhoods beyond 5 points.

NEXT REVIEW

15/01/2028

A blind spot written down is a blind spot watched.

20 questions about ACF

The most frequent questions about the framework : its scope, how it differs from regulation, and how to use it. **Framework & positioning, then "when and for whom".**

01 Does ACF replace the AI Act?

No. The AI Act sets **legal obligations**; ACF is the operational framework that helps implement them across agents. It makes you AI-Act-ready - it does not substitute for it.

02 Is ACF an official standard?

No. It is a private framework, **compatible** with NIST AI RMF, ISO 42001 and ISO 23894 : which it translates into concrete practices for agentic commerce.

03 Can ACF be used without AI?

Yes. ACF governs the **delegation of decisions**, whether to an AI agent, a script or a human. The criticality / autonomy grid applies to any delegated decision.

04 Is ACF reserved for e-commerce?

No. "Commerce" refers to any **exchange of value** that commits the organization: financial, HR, legal and operational decisions.

05 Is ACF a paid framework?

The framework and the 45 cards are **free** for teaching and non-profit use. The enterprise ecosystem (platform, support) is contractual.

06 How many agents before you start governing?

From the very first one. Risk does not come from the number of agents but from the **criticality of delegated decisions**. A single agent handling an irreversible decision is enough.

07 Do you need to be a large company?

No. ACF is **scalable**: a small business covers the bare essentials (ACF-01, 03, 13), while a large organization deploys the full cycle.

08 Who should own ACF internally?

A named executive - the **DDAO** (Delegated Decision Agent Officer) : backed by senior management. Decision sovereignty is a governance issue, not just an IT one.

09 How long does it take to get started?

A diagnostic (ACF-01) and a first mapping (ACF-02/03) fit into **half a day**. The first governed agent (ACF-08) takes a few weeks.

10 Does ACF slow down innovation?

No: it **secures deployment**. Framing things upfront prevents the kind of incident that would freeze projects for the long run.

Concepts & implementation

The second half of the questions addresses the **core notions** of the framework : sovereignty, criticality, levels · and the concrete **safeguards** in case of drift.

11 Autonomy or sovereignty?

Autonomy **can be delegated** (the agent acts); sovereignty **cannot be transferred** (humans retain the final word on critical decisions).

12 What is the criticality score?

Impact × Frequency × Irreversibility ÷ 10.

Below 15: agentifiable; 15-25: governed agent; above 25: human required.

13 Criticality (ACF-03) or risk (ACF-12)?

Criticality sets the **permitted autonomy level**; the risk matrix (Probability × Severity) assesses operational risk. Two distinct scales.

14 What is a DDAO?

Delegated Decision Agent Officer: the human officer accountable for delegated decisions , they set the frame, supervise, and hold the kill switch.

15 How many autonomy levels are there?

Four, **per decision**: Advisory, Co-decision, Supervised autonomous, Autonomous. Distinct from the 4 maturity levels (N0-N3), which characterize the agent as a whole.

16 What if an agent goes off the rails?

Ultimate human control mandates a **takeover** and a kill switch (ACF-07). Reversibility is required by design (ACF-14).

17 How do you prove who decided?

Through traceable accountability: every autonomous action is **logged, auditable, and attributable** to a responsible officer (5 pillars, ACF-14).

18 Does ACF mandate a specific technology?

No. ACF is **technology-neutral**: it applies regardless of the model, vendor, or architecture in use.

19 Is the Decision Engine required?

No. The framework and the cards are enough to govern. The **ACF Decision Engine** is an optional enterprise component, not publicly detailed.

20 Where to begin in practice?

Start with card **ACF-01** (Sovereignty Score), then follow the 6-step journey. The manual and user guide walk you through it step by step.

ACF and the existing frameworks

ACF does not reinvent compliance, it **equips it**. Every card in the toolkit concretely implements one or more requirements of the main AI governance frameworks. All 45 cards are set here against the AI Act, ISO/IEC 42001, the NIST AI RMF and the GDPR.

DECIDE · ACF-01 TO ACF-20

ACF CARD	AI ACT	ISO 42001	NIST AI RMF	GDPR
ACF-01 Sovereignty score	Art. 9	§4 · §6.1	Map	
ACF-02 Decision map	Art. 9 · 26	§6.1 · §8.2	Map	Art. 22
ACF-03 Criticality matrix	Art. 9 · 14	§6.1 · §8.2-8.3	Map · Measure	Art. 22 · 35
ACF-04 Agentic constitution	Art. 17 · 14	§5.2 · Ann. A	Govern	Art. 25
ACF-05 Agent card	Art. 11 · 13	§7.5 · A.6	Map	
ACF-06 Supervision & governance	Art. 14 · 72	§9.1	Manage · Measure	Art. 32
ACF-07 Kill switch	Art. 14	§8.1	Manage	
ACF-08 First agent briefing	Art. 11 · 9	§8 · A.6	Map · Govern	
ACF-09 Agentic decisions register	Art. 12 · 26	§7.5 · §9.1	Measure	Art. 30
ACF-10 Action plan	Art. 9	§10	Manage	
ACF-11 Governance audit	Art. 17	§9.2 · §9.3	Govern · Measure	
ACF-12 Risk assessment	Art. 9 · 27	§6.1 · §8.2-8.3	Map · Measure	Art. 35
ACF-13 Agent mandate	Art. 26 · 14	§5.3	Govern	
ACF-14 Responsibility by design	Art. 14 · 26	§5.3	Govern	Art. 37-39
ACF-15 Objective & counter-metric	Art. 9 · 15	§6.2 · §9.1	Measure	
ACF-16 Supervision test	Art. 14	§9.1 · A.6	Measure · Manage	
ACF-17 Validity review	Art. 17 · 72	§9.3 · §10	Govern · Manage	
ACF-18 Guided case study				
ACF-19 Instructor's guide				
ACF-20 Governance simulation				

Organise and Operate against the same frameworks

The Organise and Operate families do not step outside the regulatory frame, they cover the part the Decide family alone left open: **what the agent shifts inside the organisation**, and **what can be verified once it is live**.

ORGANISE · ACF-O1 TO ACF-O10

ACF CARD	AI ACT	ISO 42001	NIST AI RMF	GDPR
ACF-O1 Human-AI-agent matrix	Art. 14 · 26	§5.3 · A.3	Govern · Map	
ACF-O2 Hybrid organisation map	Art. 26	§4.1 · §5.3	Map · Govern	
ACF-O3 Critical skills	Art. 4 · 26	§7.2 · A.4	Govern	
ACF-O4 AI policy	Art. 17	§5.2 · A.2	Govern	Art. 24
ACF-O5 Shadow AI	Art. 4 · 26	§8.1 · A.9	Map · Govern	Art. 5 · 30
ACF-O6 Shadow agents	Art. 26 · 49	§8.1 · A.9	Map · Manage	Art. 30
ACF-O7 The shift of power	Art. 14 · 26	§5.3 · A.3	Govern	
ACF-O8 Exceptions register	Art. 17 · 26	§7.5 · §10	Govern · Manage	Art. 30
ACF-O9 Fleet governance	Art. 26 · 72	§8.1 · §9.1	Govern · Manage	Art. 30
ACF-O10 Affected people & redress	Art. 26 · 27 · 86	A.8	Govern · Manage	Art. 12-15 · 22

Organise and Operate against the same frameworks

The Organise and Operate families do not step outside the regulatory frame, they cover the part the Decide family alone left open: **what the agent shifts inside the organisation**, and **what can be verified once it is live**.

OPERATE · ACF-P1 TO ACF-P15

ACF CARD	AI ACT	ISO 42001	NIST AI RMF	GDPR
ACF-P1 Configuration grid	Art. 13 · 15	§8.1 · A.6	Manage	
ACF-P2 Acceptance protocol	Art. 15 · 17	§8.1 · §9.1	Measure	
ACF-P3 Ramp-up plan	Art. 15 · 26	§8.1	Manage	
ACF-P4 Supervision log	Art. 14 · 26	§9.1	Measure · Manage	
ACF-P5 Role allocation	Art. 26	§5.3 · A.3	Govern	Art. 24 · 28
ACF-P6 Escalation card	Art. 14	§8.1	Manage	Art. 22
ACF-P7 Stop test	Art. 14	§8.1 · §9.1	Manage	
ACF-P8 Incident log	Art. 26 · 73	§10.1	Manage	Art. 33
ACF-P9 Permissions review	Art. 10 · 15	A.7	Manage	Art. 5 · 25 · 32
ACF-P10 Vendor assessment grid	Art. 25 · 26	A.10	Govern · Map	Art. 28
ACF-P11 Foundation dependency map	Art. 25	§4.1 · A.10	Map	Art. 28
ACF-P12 Degraded mode & exit	Art. 15	§8.1 · A.10	Manage	Art. 28
ACF-P13 Model version log	Art. 11 · 72	§7.5 · §8.1	Manage	
ACF-P14 Holding the record	Art. 12 · 19	§7.5	Measure	Art. 5 · 30
ACF-P15 Blind spot statement	Art. 9 · 13	§6.1 · §8.2	Map	Art. 35

*Indicative mappings: they show where each card helps meet a requirement; they are not a certification. **ACF-18 to ACF-20 have all four cells empty**: they are the documents of a teaching session and put the other cards to work without implementing a requirement themselves. AI Act: Regulation (EU) 2024/1689. **Timeline**: the obligations for Annex III high-risk systems, first set at 2 August 2026, are postponed to 2 December 2027 by Regulation (EU) 2026/1744, and to 2 August 2028 for AI embedded in Annex I products; the Article 50 transparency obligations do apply from 2 August 2026. ISO/IEC 42001:2023: clauses and Annex A controls. NIST AI RMF: Govern, Map, Measure, Manage. GDPR: automated decision-making and data protection. An empty cell means no directly applicable requirement.*

Resources & references

ACF is not yet another standard: it **dialogues with existing frameworks** and makes them applicable on the ground. Here are the texts, research, and tools it draws on.

REGULATORY & NORMATIVE FRAMEWORKS

AI Act · Regulation (EU) 2024/1689 · EU

The European legal framework for AI, organized by risk level. ACF helps operationalize it.

NIST AI Risk Management Framework 1.0 · USA

Voluntary framework for AI risk management: Govern, Map, Measure, Manage.

ISO/IEC 42001:2023 · ISO

AI management system (AIMS) · certifiable governance for AI systems.

ISO/IEC 23894:2023 · ISO

Guidelines for risk management related to artificial intelligence.

OECD AI Principles · OECD

International reference framework for trustworthy AI (2019, revised 2024).

Council of Europe Framework Convention on AI · CoE 2024

First binding international treaty on AI, human rights, and the rule of law.

DOCTRINE & RESEARCH

Frank Pasquale · New Laws of Robotics · 2020

The 4 new laws of robotics: complementarity, authenticity, non-weaponization, attribution.

ISN & iDFRights work · B. Benhamou · 2026

AI, digital sovereignty, and fundamental rights in Europe.

GOING FURTHER WITH ACF

The ACF® whitepaper · acfstandard.com

The reference book to which this toolkit is the practical companion.

The manual & the 45 ACF cards

· ACF-01→20 · O1→O10 · P1→P15

Diagnostics, mapping, mandates, supervision, audit · ready to print.

acfstandard.com · online

The standard, its updates, and the enterprise ecosystem.



ACF aims to be a **bridge between these frameworks and the field**. Where the AI Act, NIST, or ISO state **what to require**, ACF equips the **how to do it** · card after card, decision after decision.

Beyond the toolkit

The toolkit provides the instruments. Three separate collections say what they are for, each in its own way. **None of the three is part of the 45 cards:** each is published and downloaded separately.

ACF Playbook

"This applies to us."

End-to-end cases, one per sector, built for projection rather than demonstration. Health, public administration, banking, insurance, industry, transport, energy, telecoms, retail. Each handbook follows the same path, from the initial situation to leaving the vendor, naming at every step the cards it uses. **The framework does not move, the context does:** exactly what one expects from a standard. The organisations described are fictional and built for learning.

COMING



www.acfstandard.com/en/acf-playbook

ACF in Action

"The problem already exists."

Eighteen real companies and schemes, analysed with the ACF grid **from public information and nothing else:** Morgan Stanley, Klarna, Air Canada, Cruise, Waymo, Workday, Replit, Cursor, DPD, Perplexity, among others. Nine teaching kits accompany the publications, each setting a reference case against an incident. The analysis never claims to know these organisations' internal arrangements.

AVAILABLE



www.acfstandard.com/en/acf-action

ACF MCP

"The framework answers inside the tool."

The standard's official Model Context Protocol server, published on npm as **acf-mcp**. It exposes the doctrine, the cards, the glossary and the regulatory mappings to any compatible assistant, Claude Desktop, Cursor, Windsurf. Twelve tools, seven of which return a structured qualification from a case: autonomy level, criticality, obligations by lifecycle phase, DDAO controls. **None of them calls a model:** the answer comes from a versioned knowledge base, and every output carries the version and the fingerprint of the doctrine that produced it. It stands as preliminary qualification, never legal advice, and always calls for human review.

AVAILABLE



acfstandard.io

Four registers, four uses. The Playbook projects, ACF in Action observes, the toolkit equips, the MCP answers. A reader who recognises their own situation in a sector handbook finds the matching cards in this manual, and in ACF in Action the evidence that the problem already arises elsewhere. One framework, seen from four distances.

The ACF tools

The **framework** is free. The **tools** help you put it into practice · from the **free diagnostic** to the **enterprise platform**.

Get started for free

FREE

ACF Score

Measure your agentic sovereignty in just a few minutes.

Free · no sign-up required

AI Act Compliance Checker

Benchmark yourself against the European AI Regulation.

Free

Flash Diagnostic

A first snapshot of your governance posture.

Free · instant results

Personalized Plan

Your ACF roadmap, generated for you.

Free

Scale up

ENTERPRISE

ACF Auditor

In-depth governance audit with documented compliance.

Upon quotation

ACF Control

Continuous oversight of your agents in production.

Upon quotation

ACF Decision Engine

The governed decision engine at the heart of your agents.

Reserved for enterprise clients

Custom deployment

Integration, DDAO (Delegated Decision Agent Officer) training, and ongoing support.

Upon quotation

The framework and the 45 cards remain free. Scan to get started : the entire ecosystem is available at acfstandard.com.



Your sovereignty score

acf-score.com · free



All tools & the enterprise offering

acfstandard.com

Delegate autonomy, retain **sovereignty**

Agents are already deciding. The real question is no longer **whether** to delegate to them, but **how far**, and under what conditions. ACF provides the operational answer: a framework for entrusting autonomy without ever giving up control.

1 Delegation is not all-or-nothing
It is calibrated **decision by decision**, according to criticality · from simple advice to autonomous action.

2 Governance is designed upstream
Mandates, thresholds and safeguards are defined **before deployment**, not after the incident.

3 Everything remains reversible
Supervision, takeover, **kill switch**: ultimate human control is never relinquished.

4 One decision, one accountable owner
Every delegated action remains **attributable** to a named human · the DDAO (Delegated Decision Agent Officer).

*"The question is no longer whether to deploy agents.
It is **how to deploy them without surrendering sovereignty.**"*

To go further: acfstandard.com



Get started · your score
acf-score.com



All ACF tools
acfstandard.com

Terms of use, trademark and intellectual property

The Handbook · 2026 Edition

Toolkit version 1.0

acfstandard.com

This document, its cards, matrices, cases, diagrams, taxonomies, decision models and the order it gives them form a work protected by copyright. The name **Agentic Commerce Framework®**, the acronym **ACF®** and the associated distinctive signs are registered trademarks. **Making the toolkit available transfers no right over those trademarks.**

© 2026 Vincent Dorange. All rights reserved.

WHAT TEACHING USE ALLOWS

The toolkit is made available free of charge to educational institutions, teachers, students and researchers, for reading and downloading, for teaching, training and non-commercial academic work, for internal workshops and case studies, and for quoting limited extracts with clear attribution.

A teacher may pass the whole file to their students, unmodified, on a closed course space or within an identified course, provided the document stays complete and keeps its version and these terms. Any use keeps the authorship, version and intellectual property notices.

WHAT REQUIRES WRITTEN PERMISSION

- reproducing, modifying, adapting, translating or redistributing a substantial part of the toolkit;
- removing, hiding or altering the copyright, trademark, version or provenance notices;
- publishing the toolkit, or a modified version, under another name, another brand or white-labelled;
- embedding the cards, matrices, cases or methods in a commercial consulting, training, certification, software or platform offering;
- selling, sublicensing or monetising the toolkit, even in part;
- extracting or reconstructing the content to build a competing product, framework or service;
- using the toolkit to train, evaluate or enrich a competing model, agent or artificial intelligence system;
- implying that an organisation, programme, course, tool or certification is approved, certified or affiliated with ACF® without written agreement.

Free access is neither an assignment of rights, nor a commercial licence, nor permission to create derivative works.

VERSION, EVIDENCE AND TRACEABILITY

This toolkit is an identified version of the ACF® corpus. Its versions and changes are documented, timestamped and kept in a traceability register. The original material is protected by copyright, and both the priority date and the integrity of each version rest on filing and timestamping mechanisms. Unauthorised distribution may give rise to the appropriate measures for the protection of those rights.

Professional use. An organisation may use the toolkit to discover the framework and evaluate it internally. Large-scale use, embedding it in operational processes, access to the digital tools, the sector editions, the certifying programmes or the deployment services fall under separate offerings. Enterprise licence, partnership, training, commercial use or sector adaptation, write through the site's contact page: acfstandard.com/en/contact

Agentic Commerce Framework® · Governing autonomy, protecting sovereignty.