GOUVERNEMENT

Liberté
Égalité
Fraternité

When AI starts to act, who is accountable?

The French Public Administration experience

Fadila Leturcq - 30th of October - NHI HR Leaders Forum 2026 - Séoul

AI entered public administration through the back door

89%

of public officials surveyed **already use AI** in their work

55%

use tools their ministries **has not approved**

Survey of ~2,000 officials in nine French administrations, published with the "Notre IA" plan, June 2026

Frame AI rather than forbid it

OCTOBER 2023

Albert

Sovereign AI built by The French Digital agency, first for frontline advisors of the governement

OCTOBER 2025

The Assistant

Built with the company Mistral AI, tested by 10,000 officials in eight ministries.



JUNE 2026

“Notre IA” plan

An integrated State strategy for AI

- **Tools**: the tools, from Albert to the AI Assistant, are now just one pillar among others
- **Skills**: training for all public officials
- **Change**: supporting managers and teams through the transformation
- **Rules**: doctrines and security rules, starting with software development
- **Steering**: factoring AI into budget decisions, backed by €655 million from the France 2030 fund

An assistant helps. An agent acts.

AI assistants are conversational tools that respond to a single request (informing, drafting, summarising) and leave it to humans to act; AI agents, by contrast, are autonomous systems capable of carrying out tasks by chaining actions together to reach a goal.

Authorship
Who is the author of an act an agent prepared or carried out?

Oversight
What does human review mean when an agent takes hundreds of steps?

Limits
Which actions can an agent take alone, and which need a person?

Trust
How do officials and citizens trust decisions an agent has shaped?

3 uses cases in the French public administration

Use case 1: Agentic AI for software development

1 · Adoption ahead of the State

Tech teams seized market agents fast. Procurement was not ready: licences often bought individually, with no view on tools, tokens or productivity.

2 · Trust built step by step

Close review at first. The public servant guides their AI agent through iterative learning cycles. The AI agent progressively improves until it can operate autonomously.

3 · Responsibility moves, it does not vanish

Developers focus on the product. The reviewer now carries the most weight. There is always a human sign-off.

THE DRAFT STATE DOCTRINE IN 10 POINTS

1. A human approves every change; agents cannot merge
2. No commercial tools where sensitive data is at stake
3. Licences bought by the State: named accounts, capped costs
4. Usage, tokens and costs reported every year
5. Agents isolated, with their own identity and minimal rights
6. Secrets never leave the trusted perimeter
7. AI output is untrusted: checked by a different tool
8. Every action logged; agents cannot alter the trail
9. No lock-in: models replaceable, code maintainable without AI
10. Same rules for contractors; keep training teams and hiring juniors

Use case 2: DiploIA, developed by the Ministry for Europe and Foreign Affairs



Is it an agent?

Technically, no: it does not plan or act on other systems yet. But it does more than assist. It produces the whole translation, a professional act an official used to perform.

The accountability question

The diplomat / civil servant may send or rely on a text in a language they cannot read. Knowledge, the first condition of accountability, is missing.

Limits worth setting

- The official stays accountable for what is sent or decided
- Use graduated by stakes: information, or rights, law and diplomacy
- High stakes: a qualified human reviews the translation
- Readers are told the text was machine-translated
- Source text and tool version are kept and traced
- Officials trained to spot what a translation can get wrong

Use case 3: Agentic AI in Education and training

01
Lesson planning
Objectives, plans, sequences

02
Content creation
Slides, exercises, videos, podcasts

03
Personalised learning
Adapted to each student

04
Grading
Suggests grades and feedback

05
Progress monitoring
Spots learning gaps

06
Continuous improvement
Adjusts the next lessons

Level	What the AI does	Teacher's role
AI-assisted	Suggests exercises and corrections	Produces and chec
AI-driven workflow	Prepares lessons, materials, draft grades	Supervises and vali
Autonomous agent	Teaches, grades and adapts alone	Prospective: major

FRANCE DRAWS A LINE

"The teacher must not delegate their power of evaluation and decision to an AI tool." — CNIL

What we learn

01	**Responsibility moves, it never disappears** It shifts from producer to reviewer, but always rests on a named person.	Software, Translation
02	**Validation must be real, not a formality** The official needs the time, the skills and the visibility to check.	Translation, Teaching
03	**Oversee the chain, not only the result** When several agents contribute, every step must be traceable.	Multi-agent
04	**Match autonomy to the stakes** The more a decision affects people's rights, the closer the human.	Translation, Teaching
05	**Decide who answers before deploying** Owner, agent identity and limits are set upfront, not after an incident.	Software, Multi-agent

Agentic AI : two warnings have been raised in 2026

APRIL 2026 · ANSSI, CYBERSECURITY AGENCY

Advises against autonomous agents in production on workstations: an agent can extend its own capabilities and bypass its permissions.

JULY 2026 · CNIL AND COUNCIL FOR AI AND DIGITAL

Civil and criminal responsibility chains remain unclear. An agent has no legal personality, yet is not a passive tool.

Closing the Responsibility Gap: The French Approach

To be accountable, you need two things

KNOWLEDGE

of what the system does
and why

CONTROL

over what it does,
including stopping it

As AI agents become more autonomous and adaptive, control becomes increasingly distributed, while accountability must remain clearly assigned. Neither the designer, the deployer nor the user fully holds both.

The French law has filled a machine-made gap before

1896 · A TUGBOAT BOILER EXPLODES

No one's fault can be proven. The Court of Cassation holds the machine's **custodian** responsible, fault or not.

2026 · AN AUTONOMOUS AGENT

The custodian rule assumes someone who directs and controls the machine. A self-directed agent challenges exactly that.

GUARDRAILS WE HAVE

- Constitutional Council, 2018 : No individual decision by an algorithm that revises its own rules without human control
- EU Court of Justice, SCHUFA, 2023 : Human intervention must be real and effective, not a rubber stamp

What we are exploring

1. Graduate uses by level of autonomy

2. Real human validation for any decision affecting a citizen

3. Full traceability of what the agent did

4. Responsibility allocated before deployment, not after

Leadership: owning accountability when agents act

1 · Judgement

Decide what to delegate, and at which level of autonomy. Refuse rubber-stamp validation.

2 · Agency

Bring shadow AI into the open with a safe space to test. Let teams co-design the rules of use.

3 · Accountability

Name who answers for each agent before it is deployed, not after the first incident.

THE LEADER'S ROLE

Leaders do not need to build agents. They must decide what agents may never do alone.

In France, every manager is to be trained on "AI and work" by the end of 2027, through a programme led by DINUM with the civil service directorate.

HR: skills, roles and careers for the agentic era

1 · Train every official

Four levels, from awareness to specialisation, on one shared platform.

2 · Redefine roles and skills

Competency frameworks per profession; new roles: reviewer, agent supervisor.

3 · Protect the talent pipeline

Keep hiring and mentoring juniors; reinvest productivity gains in skills.

4 · Count oversight as real work

Give reviewers the time and authority to check, not just to approve.

TARGETS IN THE DRAFT CIRCULAR

- 100% of managers trained on AI and work by end 2027
- AI training in every public service school from 2027
- Learning communities and internal trainers from 2026
- Impact measured on the job, 3 to 6 months later

If we stop training juniors today, who will review the agents tomorrow?

Merci