

Governing AI in small states



The duties arrive whether or not the posts do

Small administrations have divided AI oversight among more bodies than their specialist pool can staff.

Spreading market surveillance across many authorities leaves every one of them understaffed. Ireland designated 13 market surveillance authorities and four notifying authorities and published a headcount for none. What it funded was coordination. A €1.5m start-up line covers a chief executive, a case-management system, a regulatory sandbox and a technical pool for the 13 to draw on¹. Spread across those 17 designated bodies, that is under one bought specialist each, before the sandbox is built.

The pool being divided is thin where it is measured best. Digital and data staff are about 5 per cent of UK civil service headcount, against a 6 per cent benchmark across other governments and 8 to 12 per cent in regulated industry; half of digital and data recruitment campaigns failed in 2024, against 22 per cent in 2019; and technical architects sit about £30,000, or 35 per cent, below the market rate². Designating 17 bodies sends them all into that market against each other.

The central alternative is legible, which makes it both countable and cuttable. Article 126 of Poland's AI Systems Act of 3 July 2026 sets a ten-year spending ceiling of 247.79m złoty, about 45 per cent below the September 2025 proposal, and it covers a ministry unit, its systems and its overheads³. The 70 posts behind that ceiling are a planned post count recorded in the regulatory impact assessment, cut from 100 before enactment. The statute itself sets no headcount. A single visible line has to survive a budget round. Duties dispersed into thirteen existing regulators are never costed, so they are never cut and never staffed either.

The EU postponed the duties and left the capacity where it was. Regulation (EU) 2026/1744, in force since 27 July 2026, sets 2 December 2027 for stand-alone high-risk systems and 2 August 2028 for those embedded in products, and its recital 40 names the delayed establishment of national competent authorities among the reasons⁴. Treat the postponement as a funded window for creating permanent posts, and measure progress by the posts on a staffing table.

FAILED CAMPAIGNS

1 in 2

UK digital and data recruitment campaigns failed in 2024, against 22 per cent in 2019²

FUNDED SPECIALISTS

<1 per authority

Ireland's whole €1.5m start-up line spread across 13 market surveillance and four notifying authorities¹

PAY GAP

£30,000

below the market rate for a technical architect, a gap of 35 per cent²

HOW WE KNOW

The figures are read off staffing tables, an estimates brief and a national workforce review. They are not like for like: Ireland's €1.5m is one year's start-up allocation, Poland's 247.79m złoty a ten-year ceiling, the UK percentages a review of that government's own civil service. No state publishes an assessor headcount for its AI duties.

AI enters government through work on text and files that no citizen sees, then through inspection targeting, and only marginally through anything that decides about a person.

The largest deployed class everywhere that counts it is internal knowledge processing, meaning search over internal documents, transcription and anonymisation, at 124 of the 421 systems the Dutch central-government inventory could classify, ahead of inspection and enforcement at 82 of the same 421 and process optimisation at 58. Of those 421, 281 have no direct impact on citizens or businesses and 140 do⁵.

Volume is low and concentrated. Eighty-eight per cent of the 70 central-government bodies surveyed ran three systems or fewer, and 120 systems were actually in use across the whole administration out of 433 reported, with the police and the employee-insurance agency holding 33 of the live ones between them⁵. Plan for tens of systems concentrated in a handful of institutions.

Size assurance capacity to what the register actually holds. The decision-facing slice the assurance debate is built around is 47 service-delivery systems and 9 detection systems out of 433⁵, so a unit established to review decisions about people will spend its first years with little to review, while the drafting and targeting tools nobody classifies as high risk go in unexamined.

Registers undercount, and self-classification runs the other way, so a direct survey beats a register as a picture of what the state runs. Only 22 of the 433 systems bodies reported to their own auditor, 5 per cent, appeared in the national algorithm register; the same bodies rated 155 of their 287 live and in-development systems minimal risk and 30 high⁵. Ask departments what they run. The register will not tell you.

TOUCH NO CITIZEN

281 of 421

central-government AI systems classified as having no direct impact on citizens or businesses⁵

SMALL ESTATES

88 per cent

of 70 central-government bodies ran three AI systems or fewer⁵

REACHED THE REGISTER

22 of 433

systems reported to the auditor that also appeared in the national algorithm register⁵

HOW WE KNOW

A court of audit survey of 70 central-government organisations, all of which responded, covering 433 systems as at October 2024. Twelve could not be classified from their descriptions, so the function split is over 421. Everything is self-reported and the auditor performed no independent check on correctness or completeness.

Pilots die at the legal basis and the funding cliff

Government AI stops for legal and budget reasons, not because the model fails.

Six of 40 publicly funded Danish AI projects finished and went into operation. Sixteen finished without reaching operation and 18 were still running, and of the 16, four ended on the legal basis and four on performance in operation⁶.

A municipal clinical tool that predicted about 50 per cent of hospital admissions at about 80 per cent precision was stopped because no sufficiently clear legal basis for operation could be identified in national law, and 68 per cent of participating municipalities and regions reported legal challenges, second only to data at 79 per cent⁶. Authority to develop and train a system is not authority to run it.

The funding cliff is a separate failure and needs a separate fix. Pool money paid for development while operating money had to come from the receiving organisation's own budget, and several business cases omitted operating and retraining costs entirely⁶. Projects with good results stopped because their own organisation ranked them below other calls on that budget.

Systems in government stop for reasons that have nothing to do with whether the model works.

Screen the legal basis and the year-two budget before the model.

Performance is rarely the reason. Of 141 terminated Dutch central-government systems, the owning bodies judged 82, or 58 per cent, to have performed as expected or better⁵. The screening question for a candidate use case is whether there will be a lawful basis for running it, a case or content system to run it inside, and a budget line for it in year two. Ask all three before the pilot.

REACHED OPERATION

6 of 40

publicly funded Danish AI projects; of the 16 that finished without going live, four failed on the legal basis⁶

STILL PERFORMING

82 of 141

terminated Dutch central-government systems that their owners judged to have performed as expected or better⁵

LEGAL CHALLENGES

68 per cent

of participating Danish municipalities and regions reported them, second only to data at 79 per cent⁶

LIMITS

The published Danish evaluation carries a DRAFT marking, and its 18 / 16 / 6 split does not reconcile with the management summary's count of six in operation, five continued in another form and eleven still running. The figures are self-reported, were collected in spring 2024, and half the sample sat in health, where medical-device rules apply on top of everything else. The Dutch 82 of 141 is the organisations' own judgement of systems they had already abandoned, which invites post-hoc generosity.

Cost per assured decision cannot be computed for any oversight model, because every state publishes a numerator and none publishes a denominator.

Poland's central model carries 70 planned posts under a ten-year ceiling of 247.79m złoty in Article 126 of its AI Systems Act of 3 July 2026³. Spain's, after three years of existence, carries 31 established civil-service posts, 10 of them in the two technical subdirectorates and five of those systems-analyst posts⁷. Ireland's federated model carries €1.5m of start-up money and no headcount, spread across 13 market surveillance and four notifying authorities, and covering a chief executive, a case-management system, a sandbox and a technical pool¹. Latvia's 13 authorities carry no published post count at all⁸. Two of the four publish no post count, and the two that do are not measuring the same thing.

The denominator is zero. No designating state has published a completed high-impact assessment, so any cost-per-decision figure quoted today is invented, and the choice among oversight models has to be made on posts and mandate until the first assessments appear.

Build the counter first, because the duty to count already exists. Article 70(3) requires an annual assessment of whether the authority holds enough permanently available personnel with in-depth understanding of AI, data, data protection, cybersecurity and fundamental rights, and Article 70(6) requires a report to the Commission on financial and human resources every two years⁴. Publishing completed assessments beside filled technical posts in that report makes the ratio exist by the second cycle.

Two proxies are measurable now. Filled technical posts per designated authority, on which a federated model is worst by construction, since the same thin pool is divided further. And posts created per euro committed, on which a central model is at least legible, since its post count is written down.

Treat a euro figure with no staffing table behind it as unfunded. Bought capacity never reaches one, and about a third of UK civil service digital and data professionals were contractors, agency or temporary staff on autumn 2023 estimates². Ten filled technical posts is a stronger position than a million euro with nothing on the table behind it.

The numerators are not like for like, and that is itself the finding. One is a planned post count recorded in an impact assessment, one a ten-year spending ceiling written into a statute, one a staffing table of established posts, one a single year's start-up allocation. Even once a denominator exists, no two of them can be divided by it and compared.

Take internal text work first, case triage second, targeting third

An administration whose AI-competent staff fit round one table should take the three classes of use case in series, because each later wave consumes the assurance and integration capability the earlier one builds.

Wave one is internal knowledge work with no external addressee: search over internal documents, drafting and summarising correspondence, transcription, translation, code conversion, and answering staff questions on HR, procurement and budget rules. It goes first because it is already the largest deployed class, 124 of the 421 classified systems in the Dutch central-government inventory⁵, and because it is bought, which converts a specialist problem into a procurement problem. The wave is complete only when a named owner reports weekly active users, because reach is not habit.

Wave two is case triage, routing, and steering enquiries from the phone into self-service. It is the first wave whose output touches a citizen, and it fails at integration. Routing that does not write into the case-management system creates a second queue and shortens nothing. Retraining is compulsory because practice drifts. One municipality found its case categorisation had shifted over three years, had a caseworker re-code 432 cases by hand and removed 690 more as unrepresentative⁶.

Wave three is targeting for inspection and enforcement, at 82 of the 421 classified systems the second-largest class after knowledge processing⁵. It needs three things the earlier waves do not: outcome-labelled historical data whose labels reflect current law and current practice, a named inspection lead personally answerable for the selection rule, and multi-year funding, since a targeting model that stops being retrained repeats last year's suspicions.

Eligibility, entitlement and individual risk scores about people are not a fourth wave. They are a separate decision that waits on a written legal basis for production, secured before development starts. The Danish tools that cleared every technical bar and still could not run failed on that point⁶.

KNOWLEDGE PROCESSING

124 of 421

classified central-government systems, the largest single class⁵

TARGETING

82 of 421

classified systems are inspection and enforcement, the second-largest class⁵

REACHED OPERATION

6 of 40

Danish projects, all of them in areas of low legal complexity⁶

THE FOUR PRECONDITIONS

A wave opens only when all four are met. An identified data source of adequate quality. A named service owner in the receiving unit who signs for the service outcome. The case or content system the output lands inside. A committed budget line for operating and retraining in the year after the pilot money stops.

Only cases that survive Gate 2 should reach a specialist, and Gates 1 and 2 are answerable from a written use-case description by the programme owner.

Gate 1 registers. It asks whether the tool meaningfully assists, supplements or automates a decision with public effect, or interacts directly with the public⁹. A tool meeting either limb earns a public register record; a tool meeting neither, such as text recognition for archiving or an internal diary scheduler, stops at an internal inventory entry.

Gate 2 determines high impact by asking whether the output is the principal basis for a decision with legal, material or significant effect on a person¹⁰. Any case sitting in one of the AI Act's eight Annex III areas passes without further argument. Rebuttal is allowed only in writing, before deployment, on the four Article 6(3) grounds, and never where the system profiles people⁴.

Gate 3 sets depth, and Canada's Algorithmic Impact Assessment is the scoring instrument already built for it. Eight risk panels are ordered by how reversible and how lasting the effect is, with level cut points at 25, 50 and 75 per cent of the maximum raw score and a 15 per cent deduction once documented mitigations reach 80 per cent of their applicable maximum. Obligations attach by level. From Level II, one qualified peer reviewer and an individualised plain-language explanation wherever a decision denies a benefit or service; Levels III and IV add a human final decision, and Level IV two published peer reviews and sign-off rising to the head of the administration¹¹.

Do not copy Level IV. Its two published peer reviews, in a jurisdiction with a handful of qualified people, mean either reviewing a colleague's own work or paying the incumbent vendor's competitor, and its sign-off ladder ending at the head of the administration has no small-state analogue. The throughput shows which tier a thin state can run. Canada has published 39 algorithmic impact assessments in the seven years since 2019¹¹. The register tier, which asks only for the Gate 1 description, grew from 59 records in May 2025 to 143 by August 2026⁹.

Self-classification is the weak joint. A system that is not high risk escapes the strict requirements, so the incentive runs one way, and a rebuttal file drafted by the delivery team turns the triage into a paperwork tax that produces no assurance⁵.

SELF-RATED MINIMAL

155 of 287

live or in-development central-government systems⁵

SELF-RATED HIGH

30 of 287

with three bodies classifying fingerprint comparison at three different levels⁵

RISK ASSESSMENT ON FILE

132 of 287

systems held a documented assessment of any kind⁵

DO NOT DUPLICATE THE PAPERWORK

Do not run Gate 3 beside a data protection impact assessment that already covers the same ground. Since 27 July 2026, Article 27(4) lets the deployer discharge the obligation by cross-reference to the relevant sections of that assessment, or by lifting those parts across, without repeating the work⁴.

A workforce budget that names only a training line has funded one of the four things it must buy.

Tier one, literacy for everyone, is delivered by a central self-paced platform and by nothing else, because that is the only format whose unit cost does not rise with headcount. It discharges the Article 4 literacy duty⁴. It produces competent users. It does not produce assessors.

Tier two, the professional core who build and run the systems, is bought with established posts and a grow-your-own pipeline, and its unit cost is a salary: about £61,000 a year for a permanent digital or data employee against about £182,000 for the contractor doing the same work². No quantity of one-day courses produces this tier.

Tier three, leaders, is intensive facilitated contact with a bounded population, so it is expensive per head and cheap in total. A leader who has done it can commission an impact assessment and sign for it. Performing one takes a tier-four specialist.

Tier four, the assurance specialists, is bought with posts on a staffing table or with access to someone else's roster, and it is the tier small administrations publish almost nothing about. A national AI agency in its third year of existence carries 10 technical posts out of 31, five of them systems analysts⁷. Article 26(2) requires oversight by natural persons with the necessary competence and authority⁴, so training the other three tiers never satisfies this one.

The pay lever closes about a third of a gap of roughly £30,000 on a technical architect's salary, sits outside the headline pay remit, and is normally funded by recycling savings from reduced use of contractors², so bought capacity and the allowance that would make the post fillable draw on the same money. No small administration publishes a before-and-after retention series for a digital pay scheme, so the effect on retention is unmeasured.

CONTRACTOR YEAR

£182,000

against about £61,000 for the permanent employee doing the same work²

BOUGHT-IN STAFF

~1 in 3

UK civil service digital and data professionals were contractors, agency or temporary staff on autumn 2023 estimates²

PUBLISHED FOR TIER FOUR

10 of 31

posts technical on the staffing table of a national AI agency in its third year⁷

THREE INSTRUMENTS THAT NEED NO PAY CASE

Fixed-term expert fellowships, whose one design choice is whether there is a route into a permanent post. Secondment from a larger regulator. And a shared roster carrying an up-to-date overview of specialists' skills, so people can move between institutions at low administrative cost.

Who should perform high-impact review?

No designating state has published a completed high-impact assessment, so the models can be compared on published post counts and on mandate, and on nothing else. Poland plans 70 posts, Spain has 31 with ten technical, and Ireland and Latvia publish none^{1,3,7,8}.

	A • One central assessment unit	B • Each sectoral regulator	C • Borrowed assessors
BENEFITS	One pool that can be counted, queued and defended in a budget round, with one employer bidding for specialists.	No new legal person and no new vote, and domain judgement stays where the systems are.	Cost scales with cases. A patent institute shared by three states reaches 150 examiners through about seven employed functions ¹² .
TRADE-OFFS	A visible budget line loses fights that thirteen invisible fractions of existing posts survive. Poland's ceiling was cut about 45 per cent ³ .	Thirteen readings of the same boundary, with the coordination bought back separately through a €1.5m office ¹ .	Borrowed capacity assesses but cannot decide, and every case leaves the building.
RISKS	An overloaded unit converts backlog into silent approvals, since a missed statutory deadline counts as agreement.	The duty is allocated and the capacity is not, and Article 70(3) then applies thirteen times over ⁴ .	No Member State has published a shared arrangement with another, and a roster drawn from vendors buys review from the supply side.
CONDITIONS	Permanent posts on a staffing table, a stated clock with a refund when it is missed, and explicit intake rationing.	The coordinator named in law with the single point of contact attached, one funded shared technical pool, one binding boundary guide.	National law making the borrowed report admissible, plus secrecy and conflict rules that outlast the engagement.

RECOMMENDATION

Concentrate the assessors in one unit and lend them to the sectoral regulators, which keep their own decisions and enforcement powers. Create the posts on a staffing table before 2 December 2027. A single line is the only capacity a state can count and defend.

OWNER • MINISTER FOR DIGITAL GOVERNMENT DECIDE BY • Q1 2027

RESIDUAL RISK • HIDDEN QUEUE

One unit is one queue, so publish the clock, the intake rule and the cases waiting. Treat any designated authority that reports no filled AI-competent post as undesigned in fact.

How should the state obtain assurance capacity it cannot recruit?

Half of UK digital and data recruitment campaigns failed in 2024, and a bought specialist costs about three times the permanent post². The mandate stays with the central unit either way; this choice settles only where its people come from, and borrowed people can staff a unit that keeps the decision.

	A · Contract it in	B · Borrow it
BENEFITS	Capacity arrives on the contract date, with no recruitment campaign to win first, and one contract can serve thirteen bodies where thirteen hires cannot.	A joint sandbox with a neighbouring state needs only a decision under Article 57(1) ⁴ , a secondment needs a bilateral arrangement, and cost scales with cases. This is the cheapest legal entry available.
TRADE-OFFS	Money in an estimate leaves nothing on the staffing table when the year ends, and the capacity goes with the contract.	The shared unit produces the report and a national official still signs the consequence, so the state keeps one post competent enough to read that report and refuse it.
RISKS	Article 26(2) and Article 27 keep oversight with the deployer ⁴ , so without an internal reader the state has bought a file nobody in it can read; the allowance that might make the post fillable comes from the same contractor line ² .	No Member State has published a shared or joint arrangement with another, and a roster bought from the supply side reviews its own market.
CONDITIONS	Name the internal signatory before the contract is let, time-box the arrangement with a dated exit into permanent posts, and publish the staffing table beside the spend.	Make the borrowed report admissible in the domestic procedure, bind borrowed staff by secrecy and conflict-of-interest rules that survive the engagement, and keep one salaried reviewer whose job is to refuse a bought assessment.

RECOMMENDATION

Borrow first: open a joint sandbox with a neighbouring state and second assessors from a larger regulator. Buy only the residue, time-boxed to 2 December 2027 with a dated exit into permanent posts.

OWNER · MINISTER FOR DIGITAL GOVERNMENT, WITH THE FINANCE MINISTRY

DECIDE BY · Q2 2027

RESIDUAL RISK · NOTHING LEFT BEHIND

Bought capacity disappears at year end while a post persists. Report filled technical posts annually beside the spend, and cut the bought share to fund the allowance.

A contract clause binds the vendor more than a year before the statute does

On bought systems, standard AI clauses produce technical documentation, event logs and a right of inspection roughly fifteen to twenty-three months before the AI Act's high-risk duties reach the same supplier.

The clause self-executes against a named counterparty. The Commission's model clauses oblige the supplier to cooperate in an audit by or on behalf of the buyer, let the buyer publish the conclusions, and put a supplier who fails to remedy an identified defect within the stated term in default by operation of law¹³. That last provision is the whole difference. The buyer triggers the sanction alone, with no designated authority, no notified body and no market surveillance headcount involved.

Adopt the Commission's models and do not draft your own, then switch the optional articles on. Four of them carry an optional marker and are the ones a supplier with leverage strikes first: real-time access to logs, help explaining an individual decision to the person affected, access to source code, data sets and development method, and the right to publish the system's entry in a public register¹³. Each is a term the state either wins at signature or does not hold at all.

Buyer conditions have moved an AI product exactly once on the published record, and only behind a coalition. A negotiator acting for a whole sector advised its institutions against a large cloud vendor's assistant on four high data-protection risks, and by September 2025 recorded two of them mitigated by changes the vendor made to the product worldwide, under audit rights that also bind its subprocessors¹⁴. No single ministry has obtained the equivalent.

The lever reaches about half the estate. Of 287 live and in-development central-government systems in one national inventory, 83 were bought outright and 57 co-developed, so a supplier clause governs 140 of that 287 and reaches none of the 115 built entirely in-house⁵. Impose the same documentation, logging and inspection duties on internal delivery teams by policy on the day the clauses are adopted. Budget the audit as well as the right: the model allows one audit a year, with the buyer paying its own auditor and a reasonable fee to the supplier unless a breach is found¹³.

PRODUCT CHANGE

2 of 4

high risks in a buyer coalition's published assessment mitigated by the supplier within nine months¹⁴

BEYOND REACH

115 of 287

central-government systems built entirely in-house, where no supplier clause reaches⁵

CLAUSE LIBRARY

24 languages

the model clauses are free and already translated, so drafting time is the only cost left to waste¹³

WHAT THE CLAUSES ARE

The model clauses are a Commission working document, not law. The Commission states that they do not reflect an official position and that the buyer carries full responsibility for what it signs. They are a template to adapt, and nothing in them shifts liability away from the contracting authority.

How should the state exercise its buying power?

The one documented case of a buyer changing an AI product came from a negotiator acting for a whole sector¹⁴. The choice is whether to build that position or rent someone else's.

	A · One national negotiator	B · Join a larger buyer's framework
BENEFITS	The assessment is the lever, and one body can afford the four or five specialists twelve ministries can each afford none of.	Day-one access under Article 39(2) of Directive 2014/24/EU ¹⁵ to terms a larger buyer spent years building, with no institution to found, plus a free ride on product changes that are global.
TRADE-OFFS	The body's contract portfolio quietly becomes national AI policy, specification control leaves the department that knows the use case, and the documented coalition took five years to build its leverage.	Purchasing runs under the host state's law under Article 39(3) ¹⁵ , and terms are drafted by the larger members. In the EU's COVID-19 vaccine procurement seven Member States negotiated and twenty received terms they had no hand in drafting ¹⁶ .
RISKS	Volume without an assessment function buys a discount and no conditions, and a single mandatory negotiator is a single queue that blocks every ministry at once.	The free ride does not extend to rights: no audit article, no default trigger, no escalation route, and eligibility or the host's priorities can move without notice.
CONDITIONS	A sole mandate barring ministries from opening a supplier negotiation alone, a published assessment function, and live agreements with the target supplier before the first AI negotiation opens.	A legal route in, someone funded at home to read the aggregator's assessments and act on them, and acceptance that the specification cannot be changed unilaterally, which rules it out where a system's parameters must follow domestic law.

RECOMMENDATION

Designate one existing central purchasing body as the state's sole AI negotiator, with signing authority and a published assessment function, before any clause work starts. Bar ministries from opening a supplier negotiation alone in the same instruction.

OWNER · MINISTER FOR PUBLIC EXPENDITURE

DECIDE BY · Q4 2026

RESIDUAL RISK · RULE-TAKER

A buyer negotiating without legislated benchmarks ratifies supplier standards under a public letterhead. Report annually how many of the four optional articles survived into signed agreements.

Every step toward controlling compute costs a generation of capability

For a state of one to three million people, renting frontier models costs single-digit millions of euros a year for a schools or single-ministry deployment and low tens of millions a year to reach every civil-service desk at list price, against about €4m to €13m a year to own the silicon^{17,18,19,20}.

Renting keeps the state on this month's model with no refresh cycle to fund and no depreciating asset on the books. A national schools programme put a current assistant in front of about 20,000 students and 4,900 teachers for €4m in its pilot year, about €160 a participant, half of it paid by private partners¹⁷. That €160 is a programme average covering training and an application as well as seats, in a schools deployment, and only 7,700 of about 20,000 student accounts were active by March 2026¹⁷. It is not a seat price.

An allocated share is the cheapest posture that still creates a national institution, a service desk and a queue the state controls. Iceland's national service point into a shared European supercomputing centre, which brokers access to a machine in Finland, is a €9.3m project whose national half was paid in staff time and facilities²¹. What it buys is time on someone else's machine, granted periodically after peer review.

Ownership costs between about €4m and €13m a year. Norway's NOK 225m research cluster annualises at about €4m over a five-year life, the €65m national half of a €130m EU-cofunded AI factory at €13m^{19,20}. Either fixes one hardware generation, with no published refresh funding behind it.

Default to renting, and escalate only for a workload that fails a named test. The data cannot lawfully leave, or the service must keep running when the vendor stops serving it, or the capability is one no vendor will sell. Jurisdiction is then bought narrowly, at about €236,000 a year for a state-controlled copy held abroad under the state's own law, plus €1m to establish it²², scoped to the named dataset. Build owned capacity only where someone else covers most of the capital and the state can commit its own share again at the next refresh. A state that cannot name the funder of the second machine should not buy the first.

PER PARTICIPANT

€160 a year

a schools programme covering about 25,000 students and teachers for €4m in its pilot year, half of it privately funded, training and application included¹⁷

COMPUTE LADDER

14x

a €130m machine of a state's own against a €9.3m service point into a shared one, secured the same month^{20,21}

RESIDENCY FALLBACK

€236,000 a year

budgeted for a state-controlled off-site copy held abroad under the state's own law²²

HOW WE KNOW

The annualised figures are arithmetic on capital sums published by the funders. No operating, power or staffing cost is published for any of these machines, so every ownership figure here understates the true cost, and by an amount that cannot be calculated from public material.

How should the state secure compute and model access?

The three postures buy current capability, a queue the state controls, or silicon. Renting is the default.

	A · Rent hosted frontier models	B · Take an allocated share	C · Own the silicon
BENEFITS	Capability stays current, with no refresh cycle to fund. Single-digit millions a year buys a schools or single-ministry deployment, low tens of millions every civil-service desk at list price ^{17,18} .	The cheapest posture that still produces a national institution and a queue the state controls, entered through a national service point for about €9.3m, half payable in kind ²¹ .	Unconditional access, with no licence to renew and no export control to survive, and EU money covering half a machine's capital ²⁰ .
TRADE-OFFS	No control of the model, the retention policy or the depreciation date; the second contract sets the real price.	The state buys hours on a schedule others set, since allocation is periodic and peer-reviewed.	About €4m to €13m a year, from a research cluster up to the national half of a €130m machine ^{19,20} , and one fixed hardware generation.
RISKS	Access is revocable by a vendor, and by a government that is not yours.	An in-kind match claims the same scarce specialists, per-country shares are unpublished, and underuse is documented.	Private capital outbuilds the state on its own territory, and refresh funding is normally unpublished.
CONDITIONS	Use for data that may lawfully sit in the vendor's cloud; buy in-region residency for the rest, about €236,000 a year ²² .	Suits bursty research and prototyping. It cannot carry a service with an availability commitment.	Build only where a non-state funder covers most of the capital. Budget power, staff and the next machine first.

RECOMMENDATION

Rent hosted frontier models by default, buy in-region residency for the datasets that cannot leave, and enter shared compute through a national service point. Buy no silicon until the funder of the second machine is named.

OWNER · MINISTER FOR DIGITAL GOVERNMENT, WITH THE FINANCE MINISTRY

DECIDE BY · Q2 2027

RESIDUAL RISK · REVOCATION

Rented access can stop without the state's consent, so a service under a statutory availability duty needs a state-controlled fallback: models whose weights are published, run on rented capacity, plus an off-site copy.

Four signatures, three ministries, and the negotiator signs first

The order is forced by what the leverage depends on. The negotiator has to exist and hold standing volume before any clause text is tabled, the procurement steps need no new law, and the assessment steps do.

Publish a standing instruction naming one existing central purchasing body as the state's sole AI contract negotiator, with signing authority and a published assessment function, and bar ministries from opening a supplier negotiation alone. The one documented case of a buyer changing an AI product came from a negotiator acting for a whole sector¹⁴. The Minister for public expenditure signs this by Q4 2026, before any clause work starts.

In the first month after that, the same minister adopts the model clauses as a schedule to the existing standard IT purchasing terms: the light set as the default for any purchase containing AI, the high-risk set wherever the output is the principal basis for a decision with legal or material effect on a person. The same instrument sets the four articles the model marks optional as the negotiator's opening position, requires a register of which supplier struck which, and carries a budget line for one audit per agreement per year, covering both the buyer's auditor and the supplier's fee¹³. An unfunded audit right produces nothing.

By Q1 2027 the Minister for digital government creates the assessment posts in one central unit that lends its assessors to the sectoral regulators, which keep their own decisions and enforcement powers, inside the window closing on 2 December 2027. One unit produces one staffing table a budget round can count. The same signature issues the internal build standard, putting the same documentation, logging and inspection duties on in-house teams, adopts the three-gate triage so that only Gate 2 survivors reach a specialist, sets the adoption order at internal text work, then case triage, then targeting, and reports filled technical posts beside the spend under Article 70(6)^{4,7}.

By Q2 2027 the same minister, with the finance ministry, fills those posts and settles compute. Borrow before buying, through a joint sandbox with a neighbouring state under Article 57(1) and secondment from a larger regulator, with bought capacity only for the residue and a dated exit⁴. Rent hosted frontier models by default, buy in-region residency for datasets under a residency duty, enter shared compute through a national service point, and buy no silicon until the funder of the second machine is named^{17,20,21}.

Three years after the pilot, no uptake figures for the model clauses have been published, and the Commission calls them a working document that leaves the buyer carrying full responsibility¹³. Adopting them is a defensible bet, and nobody has shown it is majority practice. The negotiator's assessment function is the part of this sequence that cannot be deferred.

Policy Briefs



Short guidebooks for ministers and senior officials, published by [Your Institution]. Each issue follows one question from evidence to decision, and is reviewed by an independent panel before publication. Series editor: [Name Surname].

SOURCES CITED IN THIS BRIEF

IN ORDER OF CITATION

1 Oireachtas, <i>DETE Estimate for 2026</i> , 4 March 2026.	2 DSIT, <i>State of digital government review</i> , 2025.
3 Polish AI Systems Act, Dz.U. 2026 poz. 1003.	4 AI Act (EU) 2024/1689, consolidated 27 July 2026.
5 Algemene Rekenkamer, <i>Focus op AI</i> , 16 October 2024.	6 Digitaliseringsstyrelsen, <i>AI-signaturprojekter</i> , 2024.
7 Registro Central de Personal, <i>RPT: AESIA</i> , Spain, 2026.	8 VARAM, <i>AI Act implementation plan</i> , Latvia, 2025.
9 GDS, ATRS scope policy and records, UK, 2024 to 2026.	10 OMB, <i>Memorandum M-25-21</i> , White House, 2025.
11 <i>Directive on Automated Decision-Making</i> , Canada, 2025.	12 Nordic Patent Institute, about and organisation, 2026.
13 European Commission, <i>Model AI clauses</i> , 2025.	14 SURF, <i>DPIA on Microsoft 365 Copilot</i> , September 2025.
15 Directive 2014/24/EU, Article 39 (procurement).	16 European Court of Auditors, <i>Special report 19/2022</i> .
17 AI Leap Foundation, <i>AI Leap launch</i> , 9 March 2026.	18 Microsoft, <i>365 Copilot business pricing</i> , read 2026.
19 Sigma2, <i>Olivia</i> , Norway, 17 June 2025.	20 LRT English, <i>Lithuania's €130m AI factory</i> , 2025.
21 Reykjavík University, <i>Iceland's AI centre</i> , 2025.	22 ERR News, <i>Estonia's data embassy bill</i> , 2018.

HOW TO CITE

[Your Institution], *Governing AI in small states*, Policy Briefs No. 04, August 2026.

CORRECTIONS & UPDATES

Evidence moves. Corrections and revised editions are published at [institution.example]; the version of record carries this issue number.

ALSO IN THIS SERIES

01 Digital public infrastructure	02 Data and the public record
03 Procurement and vendor terms	04 Governing AI in small states · this issue
05 AI and the public workforce · forthcoming	06 Safety, incidents and redress · forthcoming