newnordics AI

The Nordic Way: Resilience, not retreat

What AI resilience means for **Denmark**, **Estonia**, **Finland**, **Iceland**, **Latvia**, **Lithuania**, **Norway** and **Sweden**, and the choices the region faces together.

NOVEMBER 2026

Introduction

AI is changing what people, companies and governments can do, and it is changing quickly. The most capable systems are built by a small number of companies in two countries, and they are becoming part of how public services, businesses and citizens work. That raises a practical question for every country in the region: what happens to the services and activities we depend on when AI becomes unavailable, performs badly or changes in ways we cannot manage?

This paper is a starting point for a wider conversation about that question in the eight countries of the New Nordics: Denmark, Estonia, Finland, Iceland, Latvia, Lithuania, Norway and Sweden. We hope the conversation leads to closer regional collaboration. The paper explains the main dependencies, sets out the evidence we could find, and describes the choices open to the region. Their merits depend on the service, the risk and the future in view, so the paper does not seek agreement on a single strategy.

We use one idea throughout. The futurist Sinéad Bovell argues that the debate on AI has split into camps that each see only part of the picture, while too little has been built to steer and safeguard the technology. Her answer is to hold co-existing realities: claims that seem to pull against each other and are both true.[1] The most important pair for this paper is about access. Access to the best AI systems matters, because losing it could put the region at a disadvantage in some fields. The ability to switch supplier, or to carry on without one, matters too. How much each matters depends on the task.

The paper's message is this: AI resilience means being able to benefit from AI while protecting important outcomes when systems, suppliers or circumstances change. Three ideas follow from it.

— A service can stay online while its quality, accountability or the skills behind it decline.

— Using the best AI systems and preparing alternatives go together. How much preparation a service needs depends on what failure would cost.

— Regional cooperation pays where it builds capabilities that countries would struggle to sustain on their own.

1 SINÉAD BOVELL, "WE HAVE LOST THE ABILITY TO HAVE A RATIONAL CONVERSATION ABOUT AI", SUBSTACK, 17 SEPTEMBER 2026.

Chapter 1. What AI resilience is

1.1 The definition

AI resilience is the ability of public bodies, businesses and societies to keep essential functions running, to adapt to disruption, and to recover from failure as AI systems and their dependencies change.

The definition has four parts.

1 Continuity

Can an important service keep operating during a disruption?

2 Adaptability

Can the service change models, suppliers or ways of working, and still use the best tools available?

3 Assurance and accountability

Can failures be found, explained and corrected?

4 Capability

Do people have the skills and the authority to manage change, including carrying on without AI for a time?

The four parts support each other. A service that can switch supplier is more likely to keep running. A team that can test different models can also spot when the current one starts to go wrong.

1.2 Two ways resilience fails

Resilience can fail suddenly, when access is cut, a cable is damaged or a system is attacked. It can also fail slowly, as small errors, unnoticed manipulation and lost trust build up. Gradual decline can do serious harm on its own, and it leaves a service weaker when a shock does come. We borrow this distinction from the philosopher Atoosa Kasirzadeh, who calls the second pattern accumulative risk and argues that it needs monitoring shared across sectors, not only plans for single emergencies. [2] The computer scientists Arvind Narayanan and Sayash Kapoor make a related point: societies underinvest in resilience, and gradual decay is so widespread that people get used to it. [3]

Slow failure is harder to see, because the service stays available. Its quality may decline. Staff may become less able to question its answers. Errors may keep affecting the same groups of people. Noticing this means watching outcomes, complaints and corrections, and whether staff can still do the work without AI, and not only whether the service is up.

---

2 ATOOSA KASIRZADEH, "TWO TYPES OF AI EXISTENTIAL RISK: DECISIVE AND ACCUMULATIVE", ARXIV 2401.07836 (V3, JANUARY 2025), FORTHCOMING IN *PHILOSOPHICAL STUDIES*. HER PAPER CONCERNS EXTREME RISKS; WE USE ONLY ITS DISTINCTION BETWEEN SUDDEN AND ACCUMULATIVE FAILURE.

3 ARVIND NARAYANAN AND SAYASH KAPOOR, "A BIG-TENT OR SMALL-TENT AI SAFETY MOVEMENT?", *AI AS NORMAL TECHNOLOGY* (SUBSTACK), 1 OCTOBER 2026, HTTPS://WWW.NORMALTECH.AI/P/A-BIG-TENT-OR-SMALL-TENT-AI-SAFETY.

	SOVEREIGNTY	RESILIENCE
What it asks	Who has the legitimate authority and control?	Do important services keep running, and recover, when something changes?
What it tends to focus on	Ownership, law and location	Continuity, adaptability, assurance and capability
How this paper uses it	Only when quoting others or naming a law	As the paper's main term, defined at the start of this chapter

1.3 Sovereignty and resilience

Two words are used in this debate. They overlap, but they ask different questions.

In practice the two often meet. Ott Velsberg, Estonia's former government chief data officer, says technological sovereignty means public control, the ability to audit, real alternatives and the ability to change provider, and not building everything at home, because a local solution is not automatically better or safer. [4] "Resilience" avoids that confusion and keeps attention on what has to keep working.

1.4 What the paper covers

The paper is about AI. But AI services run on cloud computing, data and network connections, and those layers often decide whether an AI service keeps working. A model that can be swapped is of little help if the data it needs is locked in another company's cloud. So the paper covers cloud, data and connections where they matter for AI, and leaves wider questions of software and digital security to others.

The paper covers public bodies, businesses and citizens. The evidence is deepest for public services, so most examples come from there.

Voices boxes give short quotes from the sources, so each view is heard in its own words. Each quote is labelled with the kind of organisation that said it.

4 OTT VELSBERG, LINKEDIN POST ON THE UNIVERSITY OF TARTU STUDY OF ESTONIA'S AI LEAP, 5 SEPTEMBER 2026.

VOICES 1

What sovereignty means in practice

“What cannot be audited, measured, or exited is not sovereignty; it is a promise.”

WORLD ECONOMIC FORUM, DATA SOVEREIGNTY IN PRACTICE, SEPTEMBER 2026, P. 13.

MULTISTAKEHOLDER FORUM.

(“Suverænitet er at kunne stå fast. Man kan godt være uafhængig og svag.“)

“Sovereignty means being able to stand firm. You can be independent and weak.” Our translation

JAN DAMSGAARD, COPENHAGEN BUSINESS SCHOOL, ALTINGET, 23 FEBRUARY 2026.

ACADEMIC.

(“Ísland ræður ekki hvaða líkön verða til, en ræður því hvort hægt er að skipta um þau.“)

“Iceland does not decide which models get built, but it does decide whether they can be replaced.” Our translation

ICELANDIC GOVERNMENT AI TASK FORCE, GERVIGREIND TIL GÓÐS, 2026, P. 67.

GOVERNMENT-APPOINTED ADVISORY GROUP.

Chapter 2. Why it matters now

2.1 AI and its impact

As the OECD puts it, AI has moved from answering to acting: systems now log in, fill in forms, write and run software, and carry out tasks for people and organisations. [5] For public services, this means AI can now carry out steps in a process, not only advise on them.

How far and how fast this will go is uncertain. Bill Gates, the Microsoft co-founder, wrote in August 2026 that AI could transform work across industries within a decade rather than over generations, and that there is currently no plan for how societies enter this era. He warns that attackers are gaining new capabilities faster than defenders can fix weaknesses, putting services such as hospitals, power grids and benefits systems at risk. [6] Even researchers who push back against the hype expect a large change. Arvind Narayanan and Sayash Kapoor of Princeton University describe AI as a "normal technology", and by normal they mean that it will be about as big as the industrial revolution. [7]

Few people have seen what the most capable systems can do. The AI researcher Andrej Karpathy, a founding member of OpenAI who joined Anthropic in May 2026, estimates that only about 5,000 people have access to the most advanced versions, which are not yet public. [8] People discussing AI may be looking at very different versions of it. So decisions about AI in public services may be made by people who have not seen what it can already do.

The events of 2026 show how the parts of resilience are tested, often several at once.

2.2 Access can change

On 12 June 2026 the United States government applied export controls that barred foreign nationals from Anthropic's two newest models, Fable 5 and Mythos 5. Anthropic could not check users' nationality in real time, so it switched both models off for every user in the world. The controls were lifted on 30 June, and Fable 5 came back worldwide on 1 July. Mythos 5, the version with fewer safeguards, was not restored to the wider group of partners outside the United States. [9]

On 16 October 2026 China's commerce and science ministries published a draft list of 23 AI technologies, including large models and agents, to be banned or restricted for export. A duty to register such exports took effect at once. Models already released as open source are outside the controls, and so is offering a model abroad through an API, an app or a website. Deploying a restricted model itself on servers abroad is covered. [10] This matters because many open models used in Europe come from China. Mozilla, the non-profit behind the Firefox browser, reports that seven of the ten most-used models on one large

---

5 OECD, *GOVERNMENT UNSTUCK: THE NEXT FRONTIER OF GOVERNMENT TRANSFORMATION*, OCTOBER 2026, BOX 2.2 (FROM P. 41), HTTPS://DOI.ORG/10.1787/464C0CFF-EN.

6 BILL GATES, "THE TURBULENT AI ERA IS HERE. THE CHOICES WE MAKE NOW ARE CRITICAL.", GATES NOTES, 26 AUGUST 2026: "THERE IS NO PLAN TO EASE THE ENTRY INTO THE AI ERA."

7 ARVIND NARAYANAN AND SAYASH KAPOOR, "A BIG-TENT OR SMALL-TENT AI SAFETY MOVEMENT?", *AI AS NORMAL TECHNOLOGY* (SUBSTACK), 1 OCTOBER 2026, HTTPS://WWW.NORMALTECH.AI/P/A-BIG-TENT-OR-SMALL-TENT-AI-SAFETY, ON THEIR POSITION "AI AS NORMAL TECHNOLOGY": THAT AI WILL "ONLY" BE AS IMPACTFUL AS THE INDUSTRIAL REVOLUTION.

8 ANDREJ KARPATHY, POST ON X, 12 OCTOBER 2026 (HTTPS://X.COM/KARPATHY/STATUS/2109361546505966046). THE FIGURES ARE HIS OWN ROUGH ESTIMATES. ON HIS MOVE TO ANTHROPIC: BLOOMBERG LAW, "OPENAI FOUNDING MEMBER ANDREJ KARPATHY TAKES ROLE AT ANTHROPIC", MAY 2026.

9 ANTHROPIC, STATEMENT ON REDEPLOYING FABLE 5, COVERING 30 JUNE AND 1 JULY 2026 (HTTPS://WWW.ANTHROPIC.COM/NEWS/REDEPLOYING-FABLE-5). MOZILLA, *THE STATE OF OPEN SOURCE AI* 2026, COUNTS 19 DAYS FROM 12 JUNE TO 1 JULY.

10 MINISTRY OF COMMERCE AND MINISTRY OF SCIENCE AND TECHNOLOGY OF CHINA, ANNOUNCEMENT NO. 48 OF 2026, 16 OCTOBER 2026 (HTTPS://WWW.MOFCOM.GOV.CN/ZCFB/DWMYGL/ART/2026/ART_61446AB4BEEE4F63973D4D0AB3097648.HTML), AND THE COMMERCE MINISTRY'S SPOKESPERSON Q&A OF THE SAME DAY (HTTPS://WWW.MOFCOM.GOV.CN/XWFB/XWFYRTH/ART/2026/ART_D7661DBF62794345A221404F82A25959.HTML).

routing platform are Chinese open models. [11] The US and Chinese moves both show governments preparing to decide who abroad may use their most capable AI models.

Access is not only on or off. AI developers now give their most capable models first to selected and vetted partners, as Anthropic does for cyber security and life sciences and OpenAI does for cyber defenders. Google announced on 28 October 2026 that Gemini 4 Carbon would first be available to selected cyber defenders, with access widened gradually. The French developer Mistral released the weights of its largest model openly on 29 October 2026, but offered vetted partners, such as Orange Cyberdefense, preview access to a specialised variant for cyber security work. These programmes are widening, but the top tier still goes to some organisations and not to others. [12]

The Tony Blair Institute for Global Change (TBI), a London policy institute, argues that failing to access and apply the best systems is one of the greatest threats to a country's sovereignty, and that for most countries building their own frontier model is neither feasible nor, for now, necessary. [13] That judgement assumes access stays open. If the US and China both narrow access to their most capable models, closed and open, the case for building capacity in the region, and for supporting open models built in Europe, grows.

2.3 Safeguards can block the defenders

Most AI models have safeguards that make them refuse some requests. That is usually useful. In security work it can sometimes block the people doing the defending.

In July 2026, according to its own incident report, Hugging Face, the largest public platform for sharing AI models, was attacked by an autonomous AI agent. When its security team asked commercial models to help analyse the attack, the models refused, because their safeguards could not tell defenders from attackers. The team ran the analysis on an open-weight model on its own servers instead. Hugging Face stressed that this was not an argument against safety measures, and advised others to have a capable model of their own ready before an incident. [14]

In October 2026 OpenCode, a company that makes an AI coding tool, had a similar experience during an attack. It got past the problem by joining OpenAI's programme for vetted defenders. [15]

Both cases touch several parts of resilience at once: continuity, because the work had to go on; adaptability, because the teams needed another route; and capability, because they had the skills to switch quickly.

2.4 A service can fail while it stays online

AI could free capacity that public services badly need. In Denmark, the government and the associations of municipalities and regions aim to free the equivalent of at least 30,000 full-time staff through AI by 2035. The Netcompany Institute, part of a Danish IT company that builds systems

---

11 MOZILLA, *THE STATE OF OPEN SOURCE AI* 2026, SEPTEMBER 2026, ON TOKEN VOLUME AT OPENROUTER.

12 ANTHROPIC, "ANTHROPIC CYBER MISSION", 8 OCTOBER 2026, AND THE LIFE SCIENCES VERIFICATION PROGRAM PAGE, 17 SEPTEMBER 2026. ON OPENAI'S PROGRAMME (DAYBREAK BLUE): OPENAI, DAYBREAK FOR FRONTLINE DEFENDERS (HTTPS://OPENAI.COM/INDEX/DAYBREAK-FOR-FRONTLINE-DEFENDERS/), AND TECHCRUNCH, 19 OCTOBER 2026. GOOGLE, "GEMINI 4 CARBON: OUR NEXT FRONTIER MODEL", 28 OCTOBER 2026 (HTTPS://BLOG.GOOGLE/INNOVATION-AND-AI/MODELS-AND-RESEARCH/GEMINI-MODELS/INTRODUCING-GEMINI-4-CARBON/), AND SUNDAR PICHAI, POST ON X, 28 OCTOBER 2026; ACCESS WIDENS THROUGH GOOGLE'S FAIRWIND PROGRAM. MISTRAL AI, "MISTRAL LARGE 4 (LE CHONK) IS NOW OPEN WEIGHT", OCTOBER 2026 (HTTPS://MISTRAL.AI/NEWS/MISTRAL-LARGE-4-OSS/): THE CYBER SECURITY VARIANT IS AVAILABLE TO DESIGN PARTNERS "SUBJECT TO RIGOROUS VETTING AND CONTRACTUAL SAFEGUARDS", WITH TESTING LIMITED TO SYSTEMS THEY OWN OR ARE AUTHORISED TO TEST; RELEASE DATE FROM CSO ONLINE, 30 OCTOBER 2026 (HTTPS://WWW.CSOONLINE.COM/ARTICLE/4242648/MISTRAL-RELEASES-LARGE-4-WEIGHTS-DESPITE-WARNING-OF-CYBERATTACK-UPLIFT.HTML).

13 TONY BLAIR, FOREWORD TO *SOVEREIGNTY IN THE AGE OF AI: STRATEGIC CHOICES, STRUCTURAL DEPENDENCIES*, TONY BLAIR INSTITUTE FOR GLOBAL CHANGE, 19 JANUARY 2026.

14 HUGGING FACE, "SECURITY INCIDENT DISCLOSURE — JULY 2026", 16 JULY 2026 (HTTPS://HUGGINGFACE.CO/BLOG/SECURITY-INCIDENT-JULY-2026).

15 TECHCRUNCH, "HACKERS HIT AI CODING TOOL OPENCODE", 19 OCTOBER 2026; AXIOS, 21 OCTOBER 2026. OPENCODE HAS NOT PUBLISHED A FORMAL INCIDENT REPORT.

for many public bodies, argues that the potential is larger. Two Danish professors question whether its US task data fit Danish welfare services. The analysis also does not estimate the cost of building the AI systems it relies on, which Netcompany says was outside its scope. [16] The more public services rely on AI, the more it matters that failures are found, explained and corrected.

Used carelessly, automation can damage the lawfulness and trust that public services depend on. In 2024 an automated, rule-based process in a Danish property-tax refund scheme handled 757 properties, and the agency later found that about 70 homeowners had been affected. The Danish Parliamentary Ombudsman found it unfortunate that the error was not caught in testing. The error came from simple automation, not advanced AI. The Ombudsman noted that automation makes a system directly dependent on the quality of the data it is given. [17]

2.5 Agents arrive

AI agents add a new pressure. The OECD notes that public law was built for a citizen who applies and an official who decides. Now an agent may apply on a citizen's behalf, many times over and at machine speed, and the official has to work out who authorised it. [18] It might also apply for something the citizen never meant to ask for. Agents raise questions of capacity, identity, authorisation and accountability at the same time.

The problem is already visible in banking. In October 2026 Axios reported that major US banks had found unauthorised AI agents inside customer accounts and could not reliably tell whether an action came from the customer or an agent. Most were moving towards giving agents read-only access. [19] If banks, which invest heavily in security, struggle to tell a customer from an agent, public services should expect the same.

Two responses are emerging in the region. Iceland's AI task force, a five-member group appointed by the Icelandic government in July 2026, proposes a controlled and traceable way for agents to use public services, starting with read-only access. [20] Estonia announced in June 2026 that it plans to give AI agents digital identities with limited powers, such as viewing data or making a payment up to a set amount, so that it is clear who acts on whose behalf and who is responsible. No timetable has been set. [21]

Standing still has a cost too. Iceland's task force says the biggest risk for government is inaction. The risk is not only that we get something wrong. It is also that we wait while everyone else's agents arrive.

---

16 NETCOMPANY INSTITUTE (THOMAS DAMSGAARD TØRSLØV, ANDREAS GOTFREDSEN AND SIDSE OVERGAARD BJÖRNSSON), AI IS THE KEY TO MAINTAINING A MODERN WELFARE STATE, 7 OCTOBER 2026, P. 2 (HTTPS://NETCOMPANY.COM/AI-IS-THE-KEY-TO-MAINTAINING-A-MODERN-WELFARE-STATE/). THE 30,000 TARGET WAS ANNOUNCED BY THE GOVERNMENT, KL AND DANSKE REGIONER ON 10 JUNE 2025: ALTINGET, 10 JUNE 2025 (HTTPS://WWW.ALTINGET.DK/KOMMUNAL/ARTIKEL/POLITISK-ENIGHED-OM-NYE-MAAL-KUNSTIG-INTELLIGENS-SKAL-FRIGOERE-MINDST-30000-AARSVAERK). THE CRITICISM IS FROM HELENE FRIIS RATNER (DTU) AND URSULA PLESNER (COPENHAGEN BUSINESS SCHOOL), QUOTED WITH NETCOMPANY'S RESPONSE IN INGENIØREN, 29 OCTOBER 2026 (HTTPS://ING. DK/ARTIKEL/NETCOMPANY-PREDICTS-AI-SAVINGS-POTENTIAL-90000-FTES-PUBLIC-SECTOR-PROFESSORS-DISMISS-ANALYSIS).

17 DANISH PARLIAMENTARY OMBUDSMAN, STATEMENT FOB 2025-13, 26 SEPTEMBER 2025, SECTION 3, AND THE AGENCY'S REPLY, SECTION 3.3.2 (HTTPS://WWW.RETSINFORMATION.DK/ELI/FOB/2025/13/PDF).

18 OECD, *GOVERNMENT UNSTUCK*, OCTOBER 2026, BOX 2.2.

19 AXIOS, "SCOOP: UNVETTED AI AGENTS ARE ACCESSING U.S. BANK ACCOUNTS", 22 OCTOBER 2026. THE BANKS' STATEMENTS ARE AS REPORTED BY AXIOS.

20 ICELANDIC GOVERNMENT AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, CH. 7. THE TASK FORCE WAS APPOINTED ON 1 JULY 2026 FOR THREE MONTHS AND MADE 69 PROPOSALS. ITS REPORT IS ADVICE TO THE GOVERNMENT, NOT GOVERNMENT POLICY. ONE MEMBER, LILJA DÖGG JÓNSDÓTTIR, IS THE CEO OF ALMANNARÓMUR, NNAI'S ICELANDIC PARTNER.

21 GOVERNMENT OF ESTONIA, "PRIME MINISTER MICHAL: ESTONIA TO BECOME FIRST COUNTRY TO CREATE DIGITAL IDENTITIES FOR AI AGENTS", 17 JUNE 2026. NO TIMETABLE HAS BEEN ANNOUNCED.

VOICES 2

Access and control

“Data centres in Texas or Shenzhen can lock out European customers overnight.”

A TRANSFORMATIVE AI STRATEGY FOR EUROPE, KIRA CENTER, SEPTEMBER 2026.

EXPERT STRATEGY PAPER; EMPHASIS ON SECURING ACCESS TO FRONTIER AI.

“Disruption is the means. The erosion of public trust is the end.”

NATIONAL CYBER RESILIENCE IN THE AGE OF AI, ESTONIAN MINISTRY OF JUSTICE AND DIGITAL AFFAIRS AND PARTNERS, JUNE 2026, P. 12.

GOVERNMENT SECURITY DOCTRINE.

Chapter 3. Where the region stands

The chapter asks which important services depend on outside capabilities, which alternatives within the eight countries, commercial or public, can carry real workloads, and how ready organisations are to switch. The mapping answers part of this: which suppliers and capabilities exist in the region.

3.1 The mapping

The New Nordics AI full-stack mapping was carried out in 2026 for New Nordics AI by the consultancy August Associates. It builds on lists of companies from the national nodes within New Nordics AI, August Associates and Dragonfly. Dragonfly's company data, with detail that is hard to find in public sources, was also used to check and improve the results.[22]

The dataset holds 249 company records with ties to the eight countries, and 248 of them fall within the layers the study covers. Each company is placed on one layer of the AI stack, from data centres to finished applications, and checked in two ways. A company is Fully New Nordic if its head office, majority ownership and contracting entity are all in the region. It is mature, in the study's commercial screen, if it has a product on sale today, customers can sign a contract with it directly, and it reaches a minimum scale for its layer, such as 10 megawatts of power for a data centre or 2 million euros of revenue for an application. The study records 96 companies as passing both tests. These are the study's assessments, not a certificate that a company can run important services independently. The study does not check what each company itself relies on, such as foreign models, chips or cloud.

These lists are not complete, so the mapping is a sample, not a census. A separate scouting exercise in Estonia found many more candidate companies, and Nordic Innovation, citing Dealroom, counts more than 980 venture-backed AI start-ups in six of the countries. [23]

22 NEW NORDICS *AI, NEW NORDICS AI FULL-STACK MAPPING*, PREPARED BY AUGUST ASSOCIATES, FINAL TEXT AND DATASET, OCTOBER 2026. COMPANY LISTS CAME FROM NEW NORDICS AI'S NATIONAL NODES, AUGUST ASSOCIATES AND DRAGONFLY, AND DRAGONFLY'S DATA WAS USED TO REFINE THE RESULTS. COMPANY DETAILS WERE DRAWN FROM PUBLIC SOURCES, NATIONAL COMPANY REGISTERS AND COMMERCIAL COMPANY DATABASES, MAINLY INVEN AND VALU8. THE DATASET COVERS SIX LAYERS OF THE AI STACK (L1 TO L6) IN ALL EIGHT COUNTRIES. POWER, TELECOMS AND CHIPS (L0) AND AI SERVICES AND CONSULTING (L7) ARE OUTSIDE ITS SCOPE; ONE RECORD PLACED IN L0 IS THEREFORE EXCLUDED. MATURITY THRESHOLDS INCLUDE AT LEAST 1,000 GPUS OR CPUS IN PRODUCTION FOR AI COMPUTING, AND REVENUE OF AT LEAST 2 MILLION EUROS OR TEN PAYING ENTERPRISE CUSTOMERS FOR APPLICATIONS AND TOOLING. THE JURISDICTION TEST ASKS WHETHER THE CONTRACTING ENTITY IS REGISTERED IN THE REGION AND HAS NO FOREIGN PARENT THAT CAN COMPEL ACCESS TO DATA; IT IS A SCREEN, NOT A LEGAL ASSURANCE. DATA LOCATION AND MODEL SOURCE ARE REPORTED WHERE PUBLIC SOURCES GIVE THEM. THE MAPPING WILL BE PUBLISHED ON THE NEW NORDICS AI WEBSITE, NEWNORDICS.AI.

23 NORDIC INNOVATION, *STATE OF THE NORDIC INNOVATION ECOSYSTEM*, AUGUST 2026, P. 65, CITING DEALROOM.

LAYER OF THE AI STACK	COMPANIES MAPPED	FULLY NEW NORDIC AND MATURE
Data centre infrastructure	30	10
AI computing (rented specialised chips)	3	2
Cloud computing	15	8
Data platforms and data providers	30	20
Model development and training	6	4
Model fine-tuning and hosting	2	0
Tooling, orchestration and model operations	13	4
Applications for any sector	44	15
Applications for one sector	105	33
Total	**248**	**96**

3.2 What the mapping shows

Four findings stand out.

The study identifies established regional offerings in infrastructure, data platforms and applications, with a narrower supplier base in computing, models and tooling. In the sample, applications make up 60% of the companies, data centres and data platforms 24%, and the layers in between 16%. A third of the 30 data centre companies and two thirds of the 30 data platform companies are both Fully New Nordic and mature, and so are 48 application companies, more than in any other layer. Health and energy are the strongest of the five sectors studied. In cloud computing, 8 of 15 companies pass both tests, but the study finds their capacity and range of services much narrower than those of the global platforms, so large or advanced workloads still depend on those platforms.

The middle of the stack is thin. The six companies classified in model development focus on particular languages or fields, such as Danish speech, Icelandic or clinical language, and they are at different stages. General-purpose models, chips, most AI computing and most tooling come from outside the region. To run a different model without relying on another foreign provider, a service needs an alternative model, computing power to run it, and tooling to connect it. The mapping identifies a narrow commercial supplier base in these layers. It

does not show whether a particular service could put together a regional fallback using commercial, public or in-house capabilities. The study notes that much tooling is built inside companies and not bought, so the count leaves out in-house and open-source work.

Location in the region is not the same as ownership. About three in five of the data centre companies in the mapping are owned in the region. The study finds foreign ownership concentrated among the largest operators, drawn by clean power, a cool climate and good connections. These are counts of companies, not of capacity. Applications run the other way: about 90% are owned in the region, partly because many are still young. Ownership can also move. Silo AI, founded in Helsinki, developed the open Poro and Viking models for Nordic languages. AMD bought it in 2024 for about 665 million dollars, [24] and the mapping now counts it as only partly New Nordic. Whether the Silo AI sale changed the region's access to skills, models or operating capacity is not known from public sources.

Few companies document keeping data in the region. Where data location matters, the dataset gives a definite regional or non-regional classification for about half of the companies. Five are recorded as keeping customer data within the eight countries. Company disclosures often describe hosting options, default locations or subcontractors, not firm commitments, so these figures are indications. The mapping records ownership, contracting arrangements and reported data locations that matter for sovereignty assessments such as those in the EU's proposed Cloud and AI Development Act (CADA). It does not assess conformity with CADA or show a tested exit for any particular service.

3.3 What the mapping does not show

The mapping lists companies and their products. It is not a list of which public services depend on which suppliers, and it does not measure market share, computing capacity or the region's total AI capability. Four limits matter.

First, a regional company may still depend on a foreign base model, chips, software or cloud beneath its own product. Public sources often do not say which model an application runs on. Second, the thresholds differ between layers, so the counts cannot be added up into one measure of capacity, and some classifications rest on judgement where public information is thin. [25] Third, the mapping counts what is sold. Public capacity can fall outside it: Finland's national computing centre CSC, which runs the LUMI supercomputer, does not count as mature, because access to it is allocated and not sold.

Fourth, the sample is broad but not complete. In a separate exercise, AI and Robotics Estonia (AIRE) identified 117 additional candidate entries, checked against an earlier working file of the mapping. Many concern defence, security and space, including drones and other autonomous systems outside the mapping's scope. Fourteen are currently marked as meeting the scale threshold, and eight also have positive ownership and contracting-entity flags. These results cannot simply be added to the mapping's totals. [26] National investigation can widen the list of candidate suppliers considerably, depending on scope, identity checks and shared assessment rules.

---

24 AMD, "AMD COMPLETES ACQUISITION OF SILO AI", PRESS RELEASE, 12 AUGUST 2024.

25 IN A FEW CASES THE RECORDED DECISION AND ITS SUPPORTING NOTE DO NOT FULLY AGREE, FOR EXAMPLE WHERE A SCALE THRESHOLD IS MARKED AS MET ALTHOUGH THE NOTE SAYS IT CANNOT BE CONFIRMED, OR AS FAILED WHERE THE CUSTOMER COUNT IS UNKNOWN. THE FIGURE OF 96 IS THE STUDY'S RECORDED RESULT.

26 AI AND ROBOTICS ESTONIA (AIRE), ESTONIAN SCOUTING FOR THE NEW NORDICS AI FULL-STACK MAPPING, OCTOBER 2026. IT HOLDS 144 ENTRIES: 27 ALREADY IN THE PROJECT'S WORKING FILE AND 117 ADDITIONAL ENTRIES, WITH SOURCES INCLUDING THE STARTUP ESTONIA DATABASE AND INDUSTRY ASSOCIATION LISTS. ESTONIAN FOUNDING IS STILL TO BE VERIFIED FOR 23 OF THE ADDITIONS, AND TWO ADDITIONS ARE THE SAME LEGAL ENTITY. FOURTEEN ADDITIONS ARE RECORDED AS MEETING THE SCALE THRESHOLD FOR THEIR LAYER, AND 98 ARE APPLICATIONS FOR ONE SECTOR. THE SCOUTING DOES NOT RECORD HEAD OFFICE, READY PRODUCT OR DIRECT CONTRACTING, SO IT DOES NOT APPLY EVERY TEST THE MAPPING USES. OWNERSHIP WAS READ FROM THE ESTONIAN BUSINESS REGISTER, WHERE AN ESTONIAN PERSONAL CODE WAS COUNTED AS AN ESTONIAN OWNER. E-RESIDENTS ALSO RECEIVE THESE CODES, SO A CODE DOES NOT SHOW WHERE THE OWNER LIVES.

3.4 What this means for the discussion

What matters is where regional offerings are established, where they are available with limitations, and where readiness is not established.

— **Established offerings:** data centres, data platforms and many applications have established regional suppliers today.

— **Available with limitations:** regional cloud, models for particular languages and fields, and rented computing power serve real workloads, but not at the scale of the global platforms.

— **Readiness not established:** whether an identified supplier could take over a specified service at the required quality, cost and recovery time, and which models regional applications depend on.

For resilience, the next question is which capabilities can support a particular service at an acceptable quality, cost and recovery time. Location, ownership and scale help identify candidates, but they do not answer that question on their own. Answering it needs a second step: linking suppliers to specific services, with evidence about how they operate.

Chapter 4. The choices

4.1 The central question

Using the best AI systems and being ready for them to fail are not opposites. A service can use a leading model and still keep a tested way to change it, correct it or carry on without it. How much preparation each service needs depends on two things: how much it gains from the best systems, and how much harm it would do if the AI failed, changed or disappeared.

The cells are examples of responses, not instructions. A less critical service can rely on one supplier with no fallback ready, as long as that is a deliberate choice that someone has written down and reviews. A critical service needs a prepared response, and any alternative must work at the quality, volume and speed the service requires.

This gives the paper its central question: **for each important service, how much does it need the best AI, and how ready must it be if that AI fails, changes or disappears?**

Readiness can take several forms: a tested second model, a way to find and correct errors, the authority to pause a feature, a reduced service, or people who can take over.

There are good reasons to want the best systems. They keep improving, and for some tasks those who cannot use them fall behind, so most strategies for AI independence still need some access to them.

There are also good reasons to keep a way out. In 2026 a government cut access to two leading models, and a model's own safeguards blocked the defenders of a company under attack.

Open models are one way to keep a way out. An open-weight model's trained weights are published, so others can run and adapt it on their own computers within the terms of its licence. Many of the leading open models, including the Chinese ones, are open-weight and not fully open source. [27] In September 2026, on the measures Mozilla tracks, the best open-weight models were roughly four months behind the best closed

27 OPEN SOURCE INITIATIVE, *THE OPEN SOURCE AI DEFINITION*, VERSION 1.0 (HTTPS://OPENSOURCE.ORG/AI/OPEN-SOURCE-AI-DEFINITION).

	LITTLE HARM IF THE AI FAILS, CHANGES OR DISAPPEARS	SERIOUS HARM IF THE AI FAILS, CHANGES OR DISAPPEARS
Gains little from the best AI	Use what works and costs the least, and accept the dependence	Choose a model you could replace or run yourself, even if it is not the best, and keep a tested fallback
Gains a lot from the best AI	Use the best, and accept the dependence by choice	Keep the best, and prepare now: a tested second option, a way to correct errors, and people who can step in

ones, and cost a fraction as much to use through an online service. [28] Results depend on the tests used. In late October, Artificial Analysis, whose earlier index Mozilla used, released a revised index built mostly on private tests, which makes it harder for developers to train their models on the questions. When it launched on 27 October, the best models whose weights have been released scored 32 against the leader's 58. Artificial Analysis reports that the gap has widened, although the revision also changed the tests, so the results cannot be compared directly with earlier versions. [29] This suggests the gap between the best closed models and openly released ones may be larger than earlier measures showed, and a fallback built on an open model may give noticeably weaker answers on hard tasks. Whether that matters depends on the service: a fallback has to meet its needs for language, accuracy, volume and cost, not match the strongest system on every task. Running a model yourself also brings its own costs.

But relying on open models alone has drawbacks. A strategy paper for Europe published through the KIRA Center, an independent Berlin think tank, calls it a risky bet, because open models trail the best closed models by several months and governments may restrict their spread, as China's draft export rules show. [30]

Two dependencies need to be kept apart. A copy of an open-weight model that an organisation already holds keeps working if an online service is withdrawn, as long as it can run the model, on its own hardware or on rented computing power, with the software and skills to do so. Running it at the volume a service needs takes computing capacity, and the region has few

28 MOZILLA, *THE STATE OF OPEN SOURCE AI* 2026, SEPTEMBER 2026. MOZILLA'S OWN ESTIMATE OF THE LAG IS ABOUT 4.4 MONTHS AND EPOCH AI'S ABOUT FOUR MONTHS; ON ONE WIDELY USED INDEX THE GAP IS TWO TO SIX POINTS, AND THE PRICE 30% TO 60% OF THE CLOSED MODEL'S.

29 ARTIFICIAL ANALYSIS, "ANNOUNCING ARTIFICIAL ANALYSIS INTELLIGENCE INDEX V5", 27 OCTOBER 2026 (HTTPS://ARTIFICIALANALYSIS.AI/ARTICLES/ARTIFICIAL-ANALYSIS-INTELLIGENCE-INDEX-V5). ABOUT 60% OF THE INDEX NOW USES PRIVATE TEST SETS OR HELD-OUT ANSWERS, KEPT PRIVATE TO AVOID CONTAMINATION. SCORES ARE FROM THE LAUNCH ON 27 OCTOBER 2026; LATER RELEASES MAY HAVE CHANGED THEM. VERSION 5 ALSO CHANGED WHICH TESTS ARE INCLUDED AND HOW THEY ARE WEIGHTED. ARTIFICIAL ANALYSIS REPORTS THE GAP BETWEEN THE LEADER AND THE BEST OPEN-WEIGHT MODEL AS 12 POINTS, AGAINST 3 ON VERSION 4.3; THAT MODEL (46) HAS ANNOUNCED BUT NOT YET RELEASED ITS WEIGHTS, SO IT IS NOT COUNTED IN THE TEXT. THE RESULTS RELY ON ARTIFICIAL ANALYSIS'S OWN TESTING, AND THE INDEX DOES NOT MEASURE PERFORMANCE IN PARTICULAR LANGUAGES OR SERVICES.

30 SCHNITZER, PRIVITERA AND OTHERS, *A TRANSFORMATIVE AI STRATEGY FOR EUROPE*, SEPTEMBER 2026 (HTTPS://TRANSFORMATIVE-AI.EU/). OBJECTIVE 1.1, "SECURE ONGOING ACCESS TO FRONTIER AI SYSTEMS". THE EXPERTS CONTRIBUTED IN PERSONAL CAPACITIES; THE KIRA CENTER IS THE CONTACT.

commercial suppliers of it. Access to future releases, updates and supporting software is a separate dependency, and a model you hold does not improve, so it falls further behind the newest systems.

Each side accepts that managing a dependence is different from removing it. The Carnegie Endowment for International Peace, a US think tank, warns against two errors, overconfidence about where you stand today and fatalism about the options you will have later. [31] For critical services, one approach between them is to keep buying access to the best systems and to prepare the fallback now. We call this **access with an exit.**

4.2 What the EU rules say

Three EU laws matter most here. The AI Act sets rules for AI systems themselves. The proposed Cloud and AI Development Act (CADA) and the Data Act cover the cloud and data that AI services run on.

CADA, proposed by the Commission on 3 June 2026 and not yet law, would set four levels of assurance for cloud and AI providers, based on where data sits, who owns the provider and who controls the supply chain. The levels do not test whether a buyer can switch provider. But the proposal asks public bodies to assess their risks to decide which level they need, and to consider whether using several providers or clouds is appropriate. [32]

The Data Act already deals with switching provider. Its EU rules have applied since 12 September 2025 and require providers of data processing services to remove obstacles to switching. Platform and software service providers must offer open interfaces and, at a minimum, export data in a commonly used, machine-readable format. Switching charges, including data-egress charges, end on 12 January 2027. [33] Whether this covers a given AI service depends on how it is built and sold, and it gives no right to a supplier's model weights. A legal right to switch also does not show that a switch works in practice. A practical exit test with real or representative data gives the strongest evidence.

4.3 A worked example: one service, six parts

An invented example, told from the municipality's point of view.

A mid-sized municipality runs a chat assistant. It answers residents' general questions about permits, school places and benefits, and it drafts replies for caseworkers handling individual cases. The two uses carry different routes to harm: wrong public guidance can affect many residents, while a wrong case draft can sway a decision that matters a great deal to one person. The service has six parts, and the municipality's position differs for each.

For each part, the last column says how much say the municipality has, using three postures from the Tony Blair Institute. [34] **Control** means the municipality decides how that part of the service runs and who operates it. **Steer** means someone else runs it, but the municipality shapes it through its contract and requirements. **Depend** means the municipality relies on someone else and has little say, by choice or because there is no alternative. None is better in itself. A sensible service mixes all three.

---

31 ANTON LEICHT AND SAM WINTER-LEVY, *AI BREAKOUT CAPACITY*, CARNEGIE ENDOWMENT FOR INTERNATIONAL PEACE, SEPTEMBER 2026, P. 15.

32 EUROPEAN COMMISSION, PROPOSAL FOR A CLOUD AND AI DEVELOPMENT ACT, COM(2026) 502 FINAL, 3 JUNE 2026: EXPLANATORY MEMORANDUM, RECITAL 65, ARTICLES 16, 18 AND 29, AND ANNEX II.

33 EUROPEAN COMMISSION, "DATA ACT EXPLAINED" (HTTPS://DIGITAL-STRATEGY.EC.EUROPA.EU/EN/FACTPAGES/DATA-ACT-EXPLAINED), ON CHAPTER VI OF THE DATA ACT. ON SCOPE: THE ACT DEFINES DATA PROCESSING SERVICES IN TECHNOLOGY-NEUTRAL TERMS, AND LAW FIRMS NOTE THAT SERVICES BUILT MAINLY ON APPLICATION LOGIC MAY FALL OUTSIDE IT (HEUKING, 4 AUGUST 2026).

34 TONY BLAIR INSTITUTE, *SOVEREIGNTY IN THE AGE OF AI*, 19 JANUARY 2026, CHAPTER 4, "A NEW FRAMEWORK FOR SOVEREIGN POSTURE: CONTROL, STEER, DEPEND". TBI'S FRAMEWORK CONCERNS NATIONAL STRATEGY; THE APPLICATION TO A SINGLE ORGANISATION IS OURS.

		IN THIS SERVICE	IF IT BECOMES UNAVAILABLE	WHAT LIMITS THE DAMAGE
Model	DEPEND *The supplier sets the terms*	A leading closed model, used through the supplier's API	Answers stop. Caseworkers draft by hand	A second model, already tested on the municipality's own questions and language
Harness	STEER *A contractor runs it to the municipality's requirements*	The software around the model: instructions, the municipality's rules, examples, links to case systems, and a log	If it is built inside one supplier's platform, with that supplier's own tools for instructions and skills, changing the model may mean rebuilding it	A harness the municipality can run itself, which works with models from several suppliers through open interfaces
Hosting	STEER *Set by the cloud contract*	A cloud service inside the EU	A cut to the service, or a change of terms, stops the assistant	Data and service that can move to another independent provider, or run locally, with an exit tested on real data
Data	CONTROL *If the municipality can access, export and restore the data itself*	Service documents, case notes and residents' messages	Locked into a closed format, or lost	Open formats, full export, lawful access, case records kept in a form staff can use without the system, and a copy held outside the supplier that has been restored at least once
Identity	DEPEND *The state runs the e-ID. The municipality can raise its needs through national cooperation, but cannot change it alone*	Residents log in through the national electronic ID	Residents cannot log in. Agents cannot be told apart from people	For residents, another way in, such as a phone line or a visit with photo ID. For agents, a separate route with credentials issued per task, which needs the national e-ID provider's support
People	CONTROL *Of staffing: the municipality organises training and cover, but can still lose key staff*	Two staff who understand the setup	They leave, retire or are overloaded, and no one else can run or switch the service	Up-to-date documentation, including the rules, instructions and test questions a new model also needs; a third trained person; and a yearly drill

Three situations show how the service can be tested.

The model is withdrawn. Suppose the assistant runs on Fable 5. It is switched off for all users on 12 June and comes back on 1 July, 19 days later. Answers stop overnight. With no preparation, the municipality has no tested alternative. It must choose a new model, may need to adapt or rebuild the harness, and must check the answers while residents wait. If it has already tested a second model, and its harness can swap models, the switch may be much faster. The answers may be somewhat worse. The municipality accepted that cost in advance. A second model is not an independent fallback if both rely on the same service that has failed.

A flood of automated requests arrives. Residents' own AI agents start to send questions and applications at machine speed. The service was built for people. Without limits and a traceable route for agents, caseworkers are overwhelmed, and real cases slow down. Resilience here means limits on how fast any one sender can submit, a separate, traceable route for agents that starts with read-only access, and a direct way for people to reach a caseworker.

The assistant gives residents wrong advice about benefits. The service keeps running, so nothing looks broken. But some residents miss an entitlement or a deadline. Resilience now means finding the error, stopping the affected use, keeping the records, correcting the advice and helping people get a remedy. That needs logs, a way for residents and staff to report problems, and someone with the authority to switch the feature off. Repeated complaints, rising correction rates or errors that fall on particular groups may show a decline before the service appears to fail.

4.4 Businesses face the same questions

Companies depend on the same few providers as public bodies, and they face the same choices. For them resilience is also about competition and value creation.

Lovable, the Stockholm company whose service turns written requests into working apps, makes heavy use of Anthropic's Claude models. According to Anthropic, Lovable reached 200 million dollars in annual recurring revenue within a year of its launch in November 2024. [35] Building on the best models gave it speed. It also brings a strategic dependence: in April 2026 Anthropic launched a product for visual design work, which commentators read as a step towards Lovable's territory, and Lovable's head of growth has named the large model providers as her main concern. [36] Public sources show that Lovable relies heavily on one supplier. They do not show whether it could switch, or how quickly. Its case shows that access to outside capability can create value fast while raising the question of what happens if the supplier becomes a competitor.

The same pattern runs through the market. Amazon, Microsoft and Google hold about 70% of the European cloud market, and European providers about 15%. [37] When a few companies supply both the models and the cloud, they can also compete with the companies that build on them.

Small businesses face the same questions with fewer means. A shop that uses one AI tool for customer messages and invoices may not have the means to host a model or

35 ANTHROPIC, CUSTOMER CASE STUDY "LOVABLE" (HTTPS://CLAUDE.COM/CUSTOMERS/LOVABLE). COMPANY ACCOUNT; THE PAGE ALSO GIVES A LATER FIGURE OF 400 MILLION DOLLARS WITHOUT A DATE.
36 ANTHROPIC LAUNCHED CLAUDE DESIGN ON 17 APRIL 2026. COMMENTARY AND THE REMARK BY LOVABLE'S HEAD OF GROWTH, ELENA VERNA (ON THE 20VC PODCAST), AS REPORTED IN "LOVABLE BUILT $400M IN ARR ON SOMEONE ELSE'S MODEL", THE STATE OF BRAND, 23 APRIL 2026.
37 SYNERGY RESEARCH GROUP, "EUROPEAN CLOUD PROVIDERS' LOCAL MARKET SHARE NOW HOLDS STEADY AT 15%", 24 JULY 2025.

run formal drills. Its safeguards can still be practical. It can choose tools built to work with models from several providers, so that changing provider does not mean starting again. It can keep a small set of its own test questions to compare models, so it knows which ones it could switch to. And it can make sure its data can be exported, and that it has a manual way to carry on for a week.

4.5 Options that strengthen resilience

Five options came up repeatedly in our sources. Each strengthens one or more of the four parts of resilience. They are candidates to examine, not recommendations.

a. Contracts with evidence and an exit

STRENGTHENS
Adaptability; assurance

WHAT IT IS
Contracts ask for assurance evidence for as long as the system is in use, including after the model changes, such as error and correction rates, complaints, and who is affected, and include exit terms: the model can be changed without rebuilding the service, data can be exported in open formats, and the exit has been tested with real data. Iceland's task force proposes four exit requirements. [38] EU procurement rules let buyers weigh quality and other criteria alongside price, as long as the criteria relate to the contract and can be checked, so these terms can count when contracts are awarded. [39]

COST
Longer procurement, possibly higher prices, staff time for exit tests

LIMITS, AND WHEN THEY MATTER
An exit test ages as services change, so it needs repeating, for example at each renewal or major model change. Most important for critical services on long contracts

b. Rehearsed fallbacks

STRENGTHENS
Continuity; capability

WHAT IT IS
Each organisation lists the functions that must keep running, chooses a reduced mode for each, and tests it. The Estonian doctrine counts a fallback only once it has been tested. [40]

COST
Staff time and test windows taken from live services

LIMITS, AND WHEN THEY MATTER
Where only the best model is good enough, the fallback may be a manual process, which may be too slow, costly or short-staffed to carry the normal workload for long

c. Shared testing and sandboxes

STRENGTHENS
Assurance; capability

WHAT IT IS
Countries or agencies share the work of testing AI systems against their own needs and languages, and open shared environments where companies and public bodies can test AI systems under controlled conditions before using them. A regulatory sandbox under the AI Act is a further step: testing supervised by a regulator, under an

38 ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, P. 67: A CONFIRMED EXIT ROUTE RUN WITH REAL DATA; DATA IN OPEN FORMATS WITH FULL EXPORT; SEPARABLE LAYERS; AND A VISIBLE SWITCHING COST IN EVERY MAJOR DECISION.
39 DIRECTIVE 2014/24/EU ON PUBLIC PROCUREMENT, ARTICLE 67: CONTRACTS ARE AWARDED ON THE MOST ECONOMICALLY ADVANTAGEOUS TENDER, WHICH MAY BE JUDGED ON THE BEST PRICE-QUALITY RATIO, AND MEMBER STATES MAY BAR PRICE-ONLY AWARDS. UNDER ARTICLE 67(3) AND (4), AWARD CRITERIA MUST BE LINKED TO THE SUBJECT MATTER OF THE CONTRACT AND ALLOW THE INFORMATION TENDERERS GIVE TO BE EFFECTIVELY VERIFIED (HTTPS://WWW.LEGISLATION.GOV.UK/EUDR/2014/24/ARTICLE/67/DATA.XHT).
40 *NATIONAL CYBER RESILIENCE IN THE AGE OF AI*, ESTONIAN MINISTRY OF JUSTICE AND DIGITAL AFFAIRS AND PARTNERS, JUNE 2026.

agreed plan. Each EU country must set up at least one. [41] Whoever decides what is "good enough" shapes every later choice

COST
Long-term funding and specialist staff

LIMITS, AND WHEN THEY MATTER
Each language and legal setting still needs its own test data, so sharing saves most on methods and tools. Renewing tests helps, but results still need checking against how systems perform in real use

d. Buying together

STRENGTHENS
Continuity; adaptability

WHAT IT IS
Neighbouring states buy computing capacity or services together. TBI argues this gives lower prices, guaranteed capacity and shared access. [42]

COST
Coordination, and a mandate across ministries and countries

LIMITS, AND WHEN THEY MATTER
Buying together can concentrate dependence on one supplier. Exit terms in the shared contract help, but they do not create independent capacity or remove suppliers that everyone shares. Works best where countries' needs and timing match

e. Fund the open components

STRENGTHENS
Adaptability; capability

WHAT IT IS
A joint fund for the open-source software that public systems rely on, which is often kept up by small volunteer communities. [43] It could also cover open training data and test sets in the region's languages

COST
Long-term funding for specialist maintenance and security work

LIMITS, AND WHEN THEY MATTER
Needs a way to choose which components matter most. The value is highest for components that many critical services share

4.6 Choices for governments

Some choices are public policy, and governments make them. Companies, universities and research institutions face versions of several of them too, such as what to build and whether to adopt open tools.

How much to build. Building more AI capacity in the region, from adapting open models to a regional frontier lab, gives more control and costs a great deal in computing power, talent and data. The KIRA strategy paper warns that a half-hearted European frontier project would leave Europe poorer and empty-handed. Carnegie proposes keeping the capacity to build quickly without building now. Iceland's task force proposes adapting open models instead of training an Icelandic model from scratch, which is the path Singapore's national language-model team also took.[44] Canada chose another route in October 2026. It announced a planned equity investment of up to 1.6 billion Canadian dollars, paid in stages, in Cohere, an existing Canadian developer of AI for businesses and governments, tied to commitments to keep its headquarters in Canada and to invest in computing capacity there. Pub-

41 REGULATION (EU) 2024/1689 (AI ACT), ARTICLE 57.
42 TONY BLAIR INSTITUTE, *SOVEREIGNTY IN THE AGE OF AI*, 19 JANUARY 2026.
43 ANNA TUMADÓTTIR, "OPENNESS & DIGITAL SOVEREIGNTY: WHO SUSTAINS THE COMMONS?", CREATIVE COMMONS, 30 SEPTEMBER 2026 (HTTPS://CREATIVECOMMONS.ORG/2026/09/30/OPENNESS-DIGITAL-SOVEREIGNTY-WHO-SUSTAINS-THE-COMMONS/).
44 LU AND OTHERS, *BUY VERSUS BUILD AN LLM: A DECISION FRAMEWORK FOR GOVERNMENTS*, ARXIV 2602.13033, FEBRUARY 2026.

lic investment of this kind can strengthen domestic capacity, but it also puts public money at risk.[45]

Trading energy and sites for access. The region could offer clean power, a cool climate and sites for data centres in return for guaranteed computing capacity or access to models. Taiwan has reserved capacity for public use in private AI computing centres, and Iceland's task force suggests negotiating taxes, a share of profit and public access to computing power in the same way.[46] The region already attracts data centres for these reasons, but more than a third of the data centre companies in the New Nordics AI full-stack mapping are owned outside the region. That counts companies, not capacity. Host communities need to see a return.

Preferring regional suppliers. A preference for regional suppliers in public procurement can build local capacity, or swap a foreign monopoly for a domestic one. Because the state is one of the largest customers in any digital economy, its requirements set terms for everyone else. Tying any preference to exit and evidence tests is one possible safeguard.

How fast to move to open alternatives. Moving public services to open tools makes switching easier, and moves the work of running and securing them to the public body or its contractors.

45 PRIME MINISTER OF CANADA, "CANADA ANNOUNCES LANDMARK INVESTMENT IN COHERE TO SCALE SECURE SOVEREIGN AI", 29 OCTOBER 2026 (HTTPS://WWW.PM.GC.CA/EN/NEWS/NEWS-RELEASES/2026/10/29/CANADA-ANNOUNCES-LANDMARK-INVESTMENT-COHERE-SCALE-SECURE-SOVEREIGN-AI). THE INVESTMENT IS MADE IN STAGES, ALONGSIDE NVIDIA, BELL AND SCHWARZ GROUP. COHERE COMMITS TO KEEP ITS CANADIAN INCORPORATION AND HEADQUARTERS AND TO INVEST AT LEAST 1.7 BILLION DOLLARS IN COMPUTING INFRASTRUCTURE IN CANADA BY 2030. THE RELEASE NAMES NO CURRENCY; AFP REPORTS THE AMOUNT AS 1.6 BILLION CANADIAN DOLLARS, AS PART OF A COHERE FUNDING ROUND TARGETING 7 BILLION CANADIAN DOLLARS ("CANADA BUYING STAKE IN AI FIRM COHERE TO BOLSTER TECH SOVEREIGNTY", 29 OCTOBER 2026, HTTPS://WWW.YAHOO.COM/NEWS/WORLD/ARTICLES/CANADA-BUYING-STAKE-AI-FIRM-194846884.HTML). THE RELEASE ALSO NOTES COHERE'S RECENT MERGER WITH GERMANY'S ALEPH ALPHA.

46 TAIPEI TIMES, 19 JUNE 2026; ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026.

VOICES 3

Building, buying and managing dependence

“The best strategy for an uncertain technological future is to preserve the capacity to shape it.”

ADAM COHEN, HEAD OF ECONOMIC POLICY, OPENAI, "WHAT IT TAKES TO LEAD IN AI", OPENAI GLOBAL AFFAIRS, 29 OCTOBER 2026.

AI DEVELOPER; WRITTEN ABOUT THE UNITED STATES AND ITS ALLIES.

“placing every chip on that bet would be geopolitical malpractice”

ANTON LEICHT AND SAM WINTER-LEVY, AI BREAKOUT CAPACITY, CARNEGIE ENDOWMENT FOR INTERNATIONAL PEACE, SEPTEMBER 2026, P. 15.

US THINK TANK; ON RELYING ONLY ON CONTINUED ACCESS FROM ONE BENIGN SUPPLIER.

Færni er sú fjárfesting sem stenst allar sviðsmyndir.

“Skills are the investment that holds up in every scenario.” Our translation

ICELANDIC GOVERNMENT AI TASK FORCE, GERVIGREIND TIL GÓÐS, 2026, P. 32.

GOVERNMENT-APPOINTED ADVISORY GROUP.

Chapter 5. Why act together, and the questions ahead

5.1 Why a regional view

Each of the eight countries has its own strengths and can run its own pilots. But many of the dependencies in this paper cross borders: the cables run under shared seas, the suppliers are the same, and so are the export controls in Washington and Beijing. A country of a few million people has little weight with a global AI company. Eight countries together have more, if they can agree on what to ask for.

Several of our sources make this case. Estonia's cyber doctrine calls pooling the only credible answer for the rarest skills, such as testing new AI systems and evaluating them for misuse. Canada's prime minister, Mark Carney, told the World Economic Forum in January 2026 that middle powers must act together, because if you are not at the table, you are on the menu. [47]

New Nordics AI is built on the same idea. In its own words: "Anchored nationally. Acting regionally. Focused on impact no single country can achieve alone."

The region has made joint commitments before: a Nordic–Baltic declaration on AI in 2018, and joint work by the Nordic prime ministers on AI-driven disinformation and democratic resilience in 2023. [48] Iceland's task force observes that the depth and usefulness of this cooperation are still untested. [49] The open question is which existing arrangements, national, Nordic, Baltic or European, could carry useful joint work.

Working together has costs. Summits and multi-year programmes move slowly compared with AI, and Nordic Innovation found that regional cooperation on innovation has often lacked an owner, funding and continuity. [50] National priorities also differ, as do views on who pays and who benefits.

The case for acting together is strongest where cooperation makes something possible that no country could do alone, where the Nordic–Baltic level fits better than the national or EU level, and where the gain is larger than the cost of coordinating. Five candidates stand out:

— **Shared pilots in the public sector.** Instead of each country running similar pilots, each could take a focus area and share its tools and findings, while still testing services in its own language and law. In Sweden, more than fifty public bodies are already building a shared assistant, Svea, on adapted open models. [51]

— **Shared work on harnesses.** Software around the models that works with many of them, so that changing model takes less work. This could mean maintaining existing open components, agreeing common interfaces or building reusable tools. The New Nordics AI full-stack mapping identifies 13 companies in the broader tooling layer, four of them both Fully New Nordic and mature. Further work would show which of them provide software that supports

47 *NATIONAL CYBER RESILIENCE IN THE AGE OF AI*, JUNE 2026; MARK CARNEY, "PRINCIPLED AND PRAGMATIC: CANADA'S PATH", DAVOS, 20 JANUARY 2026 (HTTPS://WWW.PM.GC.CA/EN/NEWS/SPEECHES/2026/01/20/PRINCIPLED-AND-PRAGMATIC-CANADAS-PATH-PRIME-MINISTER-CARNEY-ADDRESSES).

48 NORDIC COUNCIL OF MINISTERS, "AI IN THE NORDIC-BALTIC REGION", DECLARATION OF 14 MAY 2018 (HTTPS://WWW.NORDEN.ORG/EN/DECLARATION/AI-NORDIC-BALTIC-REGION); "NORDIC DIGITAL RESILIENCE", NORDIC PRIME MINISTERS, OSLO, 31 OCTOBER 2023, PUBLISHED BY THE PRIME MINISTER'S OFFICE OF ICELAND.

49 ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, P. 14.

50 NORDIC INNOVATION, *STATE OF THE NORDIC INNOVATION ECOSYSTEM*, AUGUST 2026.

51 AI SWEDEN, "SVEA'S THIRD PHASE ALMOST FULLY SUBSCRIBED" (AI.SE), WHICH GIVES 55 PARTICIPANTS IN PHASE 2 AND SAYS THE PHASE 3 TARGET OF 70 WAS EXCEEDED; VINNOVA, "PHASE 3 — A SHARED DIGITAL ASSISTANT FOR THE PUBLIC SECTOR", PROJECT 2025-04682 (DECEMBER 2025 TO OCTOBER 2026, SEK 3 MILLION).

switching models. The Tony Blair Institute argues that open models make an exit possible, and software that works with many models makes it practical. [52]

— **Shared testing of AI systems.** Testing in the region's languages, so that each country does not have to judge every new model alone.

— **Shared exit-and-evidence terms.** Common contract terms that every public buyer can use.

— **Shared reporting of AI incidents and near misses.** Reporting that helps problems that build up slowly, such as falling quality or repeated errors, to be noticed early.

5.2 The role of New Nordics AI

New Nordics AI can bring together practitioners and expert groups across the eight countries at working speed. That is hard for governments to do on their own, because each must first go through its own processes. NNAI is not a new layer of governance, and this paper does not propose to make it one. What is taken forward, and by whom, remains a decision for each government and organisation taking part.

5.3 Questions for the region

We end with seven questions. They are open, and the answers are for the region to find.

1. Which services must keep working, and keep working correctly, when AI fails, gives wrong answers or becomes unavailable?

2. Where does access to the best AI clearly make a difference, and where are other options good enough?

3. Which shared capabilities would help the region adopt AI, build competitive businesses and protect important services?

4. Where would working together cost more than it adds?

5. How should public services meet citizens' AI agents?

6. How can AI strengthen people's ability to understand, take part in and challenge decisions, instead of replacing their own judgement?

7. What evidence would change our current assessment, and who could gather it?

Annex A. Key topics

Each entry gives a short reference on one topic: what it is, why it matters here, what is reasonably established, what is uncertain or disputed, and the question it leaves.

A1. Connectivity and subsea cables

What it is. Cables on the seabed carry about 99% of intercontinental internet traffic. AI services hosted on other continents depend on them. [53]

Why it matters here. The region depends on a few routes. Iceland has three direct cable links to Europe, FARICE-1, DANICE and ÍRIS, run by Korda, a company wholly owned

52 ELIZABETH SEGER AND GUY WARD-JACKSON, *BUILD THE ROADS, NOT THE ENGINE: AN OPEN-SOURCE STRATEGY FOR MIDDLE-POWER AI SOVEREIGNTY*, TONY BLAIR INSTITUTE, SEPTEMBER 2026.
53 EUROPEAN COMMISSION AND HIGH REPRESENTATIVE, *EU ACTION PLAN ON CABLE SECURITY*, 21 FEBRUARY 2025.

by the Icelandic state. A fourth, AUÐUR, to Scotland is planned for 2030, and a separate route runs via Greenland. [54] The Baltic Sea has seen a series of incidents, including damage to a cable between Latvia and Lithuania in January 2026. [55]

Established. When two Baltic Sea cables were cut in November 2024, Internet Society Pulse measured that traffic was rerouted with a little extra delay and no meaningful loss of data. One measured case does not show how future cuts will play out. [56] Repair ships are scarce, and faster repair limits the damage.

Uncertain or disputed. How much of the damage is deliberate. The EU says the pattern in the Baltic suggests increasingly deliberate acts. But investigations sometimes find otherwise. In October 2025 Swedish prosecutors closed their investigation into the January 2025 damage to a cable between Sweden and Latvia, concluding that it was an accident. [57] One case does not settle the others.

The question it leaves. Should repair capacity be a shared regional task?

A2. Computing power and energy

What it is. Computing power, or compute, is the processing capacity used to train and run AI models. It sits in data centres and uses a lot of electricity.

Why it matters here. The region has clean power, a cool climate and established data centres. Iceland's data centres used about 3% of the country's electricity in 2025, and nearly all of the AI work in them came from abroad. [58] Capacity located here is not capacity controlled here: more than a third of the data centre companies in the New Nordics AI full-stack mapping are owned outside it. A foreign company with a data centre in Iceland is still bound by its home country's law.

Established. Data-centre power use is growing fast. In Norway, data centres used 3.3 TWh of electricity in 2025, about twice the 2023 level, and the energy regulator NVE calls them a wildcard in the power balance to 2030. That figure covers all data-centre use, not only AI. [59] Energy supply in some countries is tight: in Iceland, demand from data centres is far above the energy available.

Uncertain or disputed. Future demand, prices, and whether energy is now the limiting factor for AI progress. Whether a large regional compute commitment would pay off.

The question it leaves. How can the region use its energy and sites to gain bargaining power without losing local consent?

A3. Cloud, lock-in and procurement

What it is. Cloud services are computing and storage rented from a provider. Lock-in is how hard it is to leave. Procurement is how public bodies buy, and what they ask for.

Why it matters here. Governments are very large customers, so what they ask for shapes the market. In Iceland, government office data sits mainly in one foreign supplier's systems, and the task force treats the data layer as the real lock-in. [60]

---

54 KORDA NETWORKS, "ABOUT KORDA NETWORKS" (HTTPS://WWW.KORDA.IS/EN/ABOUT-KORDA); EFTA SURVEILLANCE AUTHORITY DECISION 019/26/COL, PARAGRAPHS 6 AND 7, ON GREENLAND CONNECT.

55 ANDRIUS SYTAS, "LATVIA PM SAYS BALTIC SEA OPTICAL CABLE HAS BEEN DAMAGED", REUTERS, 4 JANUARY 2026, REPUBLISHED BY THEPRINT ON 5 JANUARY 2026 (HTTPS://THEPRINT.IN/WORLD/LATVIA-PM-SAYS-BALTIC-SEA-OPTICAL-CABLE-HAS-BEEN-DAMAGED/2818745/). THE REPORT CITES STATEMENTS BY LATVIA'S PRIME MINISTER AND LITHUANIA'S NATIONAL CRISIS MANAGEMENT CENTRE; THE CAUSE WAS UNDER INVESTIGATION.

56 JANE RUFFINO, INTERNET SOCIETY PULSE, 22 NOVEMBER 2024; BRUCE D. JONES, *BULLETIN OF THE ATOMIC SCIENTISTS*, 13 FEBRUARY 2026.

57 SWEDISH SECURITY SERVICE, "INVESTIGATION INTO GROSS SABOTAGE HAS BEEN CLOSED", 13 OCTOBER 2025, REPORTING THE DECISION OF THE SWEDISH PROSECUTION AUTHORITY.

58 ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, PP. 57 TO 59.

59 NORWEGIAN WATER RESOURCES AND ENERGY DIRECTORATE (NVE), "ENERGIBRUK I DATASENTER", UPDATED 26 AUGUST 2026 (HTTPS://WWW.NVE.NO/ENERGI/ENERGISYSTEM/ENERGIBRUK/ENERGIBRUK-I-DATASENTER/).

60 ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, P. 67.

Established. Buying from a domestic company does not by itself give control. A foreign provider that stores data locally is still bound by its home law. The US CLOUD Act, for example, lets US authorities, through valid legal process, require providers under US jurisdiction to hand over data they control, wherever it is stored. [61] Lock-in now sits less in the model and more in the data and in what the system remembers while it works, according to Mozilla.

Uncertain or disputed. Whether one shared exit-and-evidence test could work across eight procurement systems. The WEF lists eight ways a sovereignty claim can fail, starting with "no practical exit". [62]

The question it leaves. Can the region agree on one exit-and-evidence test that every public buyer uses?

A4. Open and closed models

What it is. A closed model is used through its owner's online service. An open-weight model is published, so anyone can run and adapt it within the terms of its licence. Open-weight is not the same as open source, because training data and code may stay closed.

Why it matters here. A region that wants to be able to leave a supplier needs something to leave to. Open models also matter for small languages.

Established. On many public tests the best open models are close behind the best closed ones. On the measures Mozilla tracked in September 2026, they were roughly four months behind, at a fraction of the price of use. [63] A revised index published in late October 2026, built mostly on private tests, reports a substantial gap on its new mix of tasks. [64] A list price is not the full cost: running an open model needs hardware, staff and security work.

Uncertain or disputed. Whether open models are a real backup. On one software benchmark, the best open model and the best closed model tend to fail on the same tasks (a correlation of 0.72, measured by Together AI). That measures shared weaknesses and says nothing about outages. A fallback can keep easier tasks running during an outage even if it shares the leader's weak spots. Models already released remain usable by those who hold them. Future releases, updates and support are less certain, now that China has drafted export controls on AI technology.

The question it leaves. Who sustains the open models and components the region depends on?

A5. Language technology and testing

What it is. Language technology covers the models and tools that work in a given language. Testing, or evaluation, decides whether a model is good enough for a public service.

Why it matters here. The region has many languages, and the largest models are trained mostly on English. Iceland's task force warns that fluent Icelandic can hide wrong or misleading information about Iceland. [65]

Established. Iceland's task force proposes that the state own the recipe (data, meth-

---

61 CLARIFYING LAWFUL OVERSEAS USE OF DATA (CLOUD) ACT, 2018, 18 U.S.C. § 2713 (HTTPS://WWW.LAW.CORNELL.EDU/USCODE/TEXT/18/2713): PROVIDERS MUST DISCLOSE DATA IN THEIR "POSSESSION, CUSTODY, OR CONTROL", WHETHER IT IS LOCATED INSIDE OR OUTSIDE THE UNITED STATES.

62 WORLD ECONOMIC FORUM, *DATA SOVEREIGNTY IN PRACTICE*, SEPTEMBER 2026, P. 11.

63 MOZILLA, *THE STATE OF OPEN SOURCE AI 2026*, SEPTEMBER 2026. THE PAGE GIVES DIFFERENT FIGURES FOR THE SAME COMPARISON: A HEADING SAYS THE BEST OPEN MODEL TRAILS BY THREE POINTS AT 60% OF THE PRICE, WHILE A CHART ON THE SAME PAGE SHOWS TWO POINTS AT ABOUT 30%. TOGETHER AI, DEEPSWE BENCHMARK, 24 JULY 2026, AS REPORTED BY MOZILLA.

64 ARTIFICIAL ANALYSIS, INTELLIGENCE INDEX V5, LAUNCH SCORES OF 27 OCTOBER 2026: THE BEST MODELS WITH RELEASED WEIGHTS SCORED 32 AGAINST THE LEADER'S 58. THE REVISION CHANGED THE TESTS, SO IT CANNOT BE COMPARED DIRECTLY WITH EARLIER VERSIONS.

65 ICELANDIC AI TASK FORCE, *GERVIGREIND TIL GÓÐS*, 2026, P. 63.

ods and tests) and adapt open models, not train its own from scratch. The Institute of the Estonian Language runs a public leaderboard that tests models in Estonian. In June 2026 it reported that top models generally held up, while cheaper and open models were weaker, especially in Russian, and noted that open models are the only option for many institutions that cannot send sensitive data to cloud services. [66]

Uncertain or disputed. How good adapted open models are for real public uses, and who will fund testing over time. Work is under way across the region. AI Sweden released GPT-SW3 for the Nordic languages in 2023. Norway's National Library and the national computing provider Sigma2 have trained the Borealis models on national and European computers. Danish Foundation Models builds Danish models together with Danish tests. In Finland, the University of Turku and Silo AI released the Viking models in 2024. [67] The New Nordics AI full-stack mapping identifies six companies in model development, focused on particular languages or fields and at different stages. The sources reviewed here do not show whether they are ready for particular public-service uses.

The question it leaves. Should the region test AI systems together, and who would own that work?

A6. Cyber security and critical infrastructure

What it is. Cyber security protects systems from attack. Critical infrastructure covers the functions a society cannot do without, such as power, health care, payments, identity and public services.

Why it matters here. Estonia's 2026 cyber doctrine says to plan cyber defence the way the Nordic countries plan total defence. In 2017 a flaw in the Estonian ID card led to about 750,000 cards being suspended, and the country kept working because people had more than one way to prove who they were. [68]

Established. AI has made attacks cheaper. The Estonian doctrine says most nations now face more attacks than they can defend against. The Hugging Face and OpenCode cases show that safeguards on frontier models can also block defenders.

Uncertain or disputed. Who should control access to the strongest defensive tools. One view, voiced by David Sacks, a US government adviser, is to make powerful models widely available to defenders. [69] Another, taken by Anthropic and OpenAI, is to widen access gradually through vetted programmes.

The question it leaves. How can a democracy defend trust without destroying the conditions it depends on?

A7. Skills, talent and capital

What it is. Three kinds of skills matter: specialists who build and adapt AI systems, know-how inside organisations that use AI, and general AI literacy.

Why it matters here. Models and suppliers change, and skills stay. Iceland's task force calls skills the investment that holds up in every scenario. The region is good at

66 INSTITUTE OF THE ESTONIAN LANGUAGE (EKI), "EKI VÄRSKE UURING: PALJUSID TEHISARU MUDELEID ON LIHTNE PROPAGANDAGA MÕJUTADA", 4 JUNE 2026 (HTTPS://EKI.EE/UUDIS/EKI-VARSKE-UURING-PALJUSID-TEHISARU-MUDELEID-ON-LIHTNE-PROPAGANDAGA-MOJUTADA/); PUBLIC LEADERBOARD: HTTPS://MOODUPUU.EKI.EE/.

67 AI SWEDEN, GPT-SW3 RELEASE, 16 NOVEMBER 2023; SIGMA2, ON THE BOREALIS MODELS, 12 AUGUST 2026; DANISH FOUNDATION MODELS, A PROJECT OF AARHUS UNIVERSITY, THE UNIVERSITY OF SOUTHERN DENMARK, THE UNIVERSITY OF COPENHAGEN AND THE ALEXANDRA INSTITUTE (2024 TO 2027); SILO AI AND THE UNIVERSITY OF TURKU, VIKING RELEASE, 4 APRIL 2024.

68 *NATIONAL CYBER RESILIENCE IN THE AGE OF AI*, ESTONIAN MINISTRY OF JUSTICE AND DIGITAL AFFAIRS AND PARTNERS, JUNE 2026, P. 22.

69 DAVID SACKS, POST ON X, 18 OCTOBER 2026 (HTTPS://X.COM/DAVIDSACKS/STATUS/2111867605388493282). HIS ADVOCACY POSITION.

starting AI companies: Nordic Innovation, citing Dealroom, counts more than 980 venture-backed AI start-ups in six of the countries. [70]

Established. Foreign capital often moves ownership, headquarters and value out of the region as companies grow, according to Nordic Innovation. Carnegie warns that a strategy built only on education produces talent with nowhere to land at home.

Uncertain or disputed. Whether public bodies have the know-how to direct and supervise AI. Iceland's task force says staff are ahead of the official frameworks, and that time saved is not the same as productivity. In Finland, Sitra, Aalto University and the Swedish company Lovable ran a 2026 summer programme in which ten students built prototypes for eight public organisations. Its playbook reports ten user-tested prototypes, most of them proofs of concept, held back mainly by missing access to real data and by how ready the receiving organisations were. [71]

The question it leaves. Where will the region's AI talent land, and how will organisations build know-how that survives a change of supplier?

A8. Regulation

What it is. Four EU laws matter most. The AI Act sets risk-based rules for AI systems. The proposed CADA sets assurance levels for cloud and AI providers. The Data Act sets rules on switching between cloud services. The GDPR protects personal data. [72]

CADA's four proposed levels, in simplified form:

LEVEL	WHAT IT REQUIRES
1	Data is processed and stored in the Union
2	Providers show independence from third countries and openness about their software supply chain
3	Providers are owned and controlled from the EU, or from a third country the EU recognises
4	Providers fully control their supply chain, with no interference from a third country

Why it matters here. Six of the eight countries are EU members. Iceland and Norway take part in the single market through the EEA but do not vote on EU law. In October 2026, EFTA's records listed both the AI Act and the Data Act as still under scrutiny for incorporation into the EEA Agreement. [73]

Established. Trust is measurable. In OECD surveys only about four in ten people are confident that government use of AI will improve services, and about a third are confident it will protect personal information. [74] These are figures across the countries the OECD surveyed, not for this region.

70 NORDIC INNOVATION, *STATE OF THE NORDIC INNOVATION ECOSYSTEM*, AUGUST 2026, P. 65, CITING DEALROOM; LEICHT AND WINTER-LEVY, AI BREAKOUT CAPACITY, CARNEGIE ENDOWMENT, SEPTEMBER 2026, PP. 9 TO 10.

71 SITRA, "5 THINGS WE LEARNED FROM AI SUMMER PROGRAM BY AALTO, LOVABLE AND SITRA", 20 AUGUST 2026; AALTO UNIVERSITY, *AI SUMMER PROGRAM PLAYBOOK* 2026 (65 PP., CC BY 4.0), HTTPS://AIPLAYBOOK.AALTO-FI.LOVABLE.APP/.

72 REGULATION (EU) 2024/1689 (AI ACT), ARTICLE 1; REGULATION (EU) 2023/2854 (DATA ACT), CHAPTER VI; REGULATION (EU) 2016/679 (GDPR), ARTICLE 1.

73 EFTA, EEA-LEX RECORDS 32024R1689 (AI ACT, HTTPS://WWW.EFTA.INT/EEA-LEX/32024R1689) AND 32023R2854 (DATA ACT, HTTPS://WWW.EFTA.INT/EEA-LEX/32023R2854), CHECKED 4 NOVEMBER 2026; ÓTTAR KOLBEINSSON PROPPÉ, STAÐA GERVIGREINDAR, 5 OCTOBER 2026.

74 OECD, *GOVERNMENT UNSTUCK*, OCTOBER 2026, P. 38 (43% FOR IMPROVED SERVICES; 32% FOR PROTECTION OF PERSONAL INFORMATION).

Uncertain or disputed. How rules affect adoption. They can slow it, and they can build the trust it needs, depending on the rule, how it is applied and what it is applied to. Keegan McBride of TBI called the Commission's June 2026 sovereignty package a signal of weakness. [75] The EU's data protection supervisors, the European Data Protection Board and the European Data Protection Supervisor, take a rights-based view. In February 2026 they supported the aim of simpler digital rules, but urged lawmakers not to change the definition of personal data in the Commission's Digital Omnibus proposal. They also saw no need for a new rule allowing AI developers to rely on legitimate interest, because the existing GDPR can already permit this, subject to its conditions. [76]

The question it leaves. How can rules protect people and build trust without slowing useful uses of AI?

Annex B. Terms used in this paper

Agentic AI. AI that acts as well as answers. It can use tools, log in and complete tasks for a person or an organisation.

API. Application programming interface: the way software connects to a service. Most organisations use commercial AI models through the provider's API, so the model runs on the provider's systems.

Compute. The processing power used to train and run AI models.

Control, Steer, Depend. Three postures towards each part of an AI system: have the authority to decide how it runs and who operates it, shape it through what you buy, or rely on others by choice. None is better by default. (Tony Blair Institute.)

Exit-ability. The ability to leave a supplier or model and keep a service running. It counts only once an exit has been tested with real data.

Frontier model. One of the most capable AI systems available at a given time, built by a handful of companies.

Harness. The software around a model. It gives the model tools, memory and instructions. If the harness is portable, a buyer can change the model and keep the rest of the service.

Open-source model. A model that comes with the freedom to use, study, change and share it, and with its code, weights and enough information about its training data to build a similar system. (Open Source Initiative.)

Open-weight model. A model whose trained weights are published, so others can run and adapt it. Its code and training details may stay closed.

The New Nordics. The eight countries in this paper: Denmark, Estonia, Finland, Iceland, Latvia, Lithuania, Norway and Sweden.

---

75 KEEGAN MCBRIDE, LINKEDIN POST BEGINNING "TODAY, THE EUROPEAN COMMISSION RELEASED ITS NEW TECHNOLOGY SOVEREIGNTY PACKAGE", JUNE 2026 (HTTPS://WWW.LINKEDIN.COM/POSTS/KEEGANMCBRIDE_TODAY-THE-EUROPEAN-COMMISSION-RELEASED-ITS-ACTIVITY-7468010978568376321-WYFA).

76 EUROPEAN DATA PROTECTION BOARD AND EUROPEAN DATA PROTECTION SUPERVISOR, JOINT OPINION 2/2026 ON THE DIGITAL OMNIBUS PROPOSAL, ADOPTED 10 FEBRUARY 2026, EXECUTIVE SUMMARY AND PARAGRAPHS 21 AND 39 (HTTPS://WWW.EDPB.EUROPA.EU/SYSTEM/FILES/DOCUMENTS/2026-02/EDPB_EDPS_JOINTOPINION_202602_DIGITALOMNIBUS_EN.PDF). THEIR FORMAL OPINION AS EU DATA PROTECTION SUPERVISORS.

Annex C. References

How this paper was prepared.

The paper was written for the New Nordics High-Level Forum in Copenhagen on 6 November 2026. We reviewed a selection of reports, studies and articles from governments, think tanks, security agencies, researchers, companies and open-source groups. We say who is speaking in each case, and label companies, advocacy groups and AI developers as such. Running text paraphrases, and short verbatim quotes appear in the Voices boxes. Translations from Icelandic and Danish are ours. AI changes fast, so the paper covers developments up to 30 October 2026. Sources were last checked on 4 November 2026. Individual datasets and sources have their own dates.

The region's AI companies.

New Nordics AI, New Nordics AI full-stack mapping, prepared by August Associates (October 2026), with company lists from New Nordics AI's national nodes, August Associates and Dragonfly. AI and Robotics Estonia, Estonian scouting for the mapping (October 2026). The mapping will be published on the New Nordics AI website, newnordics.ai.

AI access, sovereignty and resilience.

Tony Blair Institute for Global Change, Sovereignty in the Age of AI: Strategic Choices, Structural Dependencies (19 January 2026); Seger and Ward-Jackson, Build the Roads, Not the Engine (September 2026); Seger and Ward, AI-Assurance Ecosystems (September 2026).

Leicht and Winter-Levy, AI Breakout Capacity, Carnegie Endowment for International Peace (September 2026).

Schnitzer, Privitera and others, A Transformative AI Strategy for Europe, KIRA Center (September 2026).

Mozilla, The State of Open Source AI 2026 (September 2026).

World Economic Forum, Data Sovereignty in Practice (September 2026; https://reports.weforum.org/docs/WEF_Data_Sovereignty_in_Practice_2026.pdf).

Glasze, Cattaruzza, Dammann and Douzet (eds.), The Elgar Companion to the Geopolitics of Digital Sovereignty (2026).

Jan Damsgaard, Digital Sovereignty (2026) and "Professor: Det offentlige skal ikke være værdipolitisk spydspids i kampen mod big tech", Altinget Digital (23 February 2026; https://www.altinget.dk/digital/artikel/forsker-det-offentlige-skal-ikke-foere-digitalt-korstog-det-skal-bygge-kapabilitet).

Lu and others, Buy versus Build an LLM: A Decision Framework for Governments, arXiv 2602.13033 (February 2026).

Atoosa Kasirzadeh, "Two Types of AI Existential Risk: Decisive and Accumulative", arXiv 2401.07836 (v3, January 2025), forthcoming in Philosophical Studies.

Open Source Initiative, The Open Source AI Definition, version 1.0. Heuking, "Data processing services under the Data Act" (4 August 2026). Artificial Analysis, "Announcing Artificial Analysis Intelligence Index v5" (27 October 2026; https://artificialanalysis.ai/articles/artificial-analysis-intelligence-index-v5).

Anna Tumadóttir, "Openness & Digital Sovereignty: Who Sustains the Commons?", Creative Commons (30 September 2026; https://creativecommons.org/2026/09/30/openness-digital-sovereignty-who-sustains-the-commons/).

Nordic and Baltic documents.

Icelandic government AI task force, Gervigreind til góðs (2026; official report and web edition: https://www.stjornarradid.is/verkefni/upplysingataeknimal-rikisins/gervigreind/gervigreind-til-gods/).

Estonian Ministry of Justice and Digital Affairs and partners, National Cyber Resilience in the Age of AI (June 2026).

Atli Þór Jóhannsson, "Hvað er gervigreindarfullveldi?", Staða gervigreindar (5 October 2026).

Óttar Kolbeinsson Proppé, on the AI Act in Iceland, Staða gervigreindar (5 October 2026).

Nordic Innovation, State of the Nordic Innovation Ecosystem (August 2026).

Nordic Ethical AI Expert Group, Policy Recommendations (June 2024).

Sitra, Aalto University and Lovable, AI Summer Program press release (15 May 2026).

Netcompany Institute, AI is the key to maintaining a modern welfare state (7 October 2026).

Danish Parliamentary Ombudsman, statement FOB 2025-13 (26 September 2025).

Government of Estonia, "Prime Minister Michal: Estonia to become first country to create digital identities for AI agents" (17 June 2026).

AI Sweden, "Svea's third phase almost fully subscribed"; Vinnova, project 2025-04682, "Phase 3 — A shared digital assistant for the public sector" (December 2025 to October 2026).

Nordic Council of Ministers, declaration on AI in the Nordic-Baltic region (14 May 2018). Nordic prime ministers, "Nordic digital resilience" (Oslo, 31 October 2023).

Korda Networks, "About Korda Networks". Norwegian Water Resources and Energy Directorate, "Energibruk i datasenter" (updated 26 August 2026).

Swedish Security Service, "Investigation into gross sabotage has been closed" (13 October 2025).

AI Sweden, GPT-SW3 release (16 November 2023); Sigma2 on the Borealis models (12 August 2026); Danish Foundation Models (Aarhus University, University of Southern Denmark, University of Copenhagen and Alexandra Institute); Silo AI and the University of Turku, Viking release (4 April 2024).

Sitra, "5 Things We Learned From AI Summer Program by Aalto, Lovable and Sitra" (20 August 2026).

Aalto University, AI Summer Program Playbook 2026 (https://aiplaybook.aalto-fi.lovable.app/).

Institute of the Estonian Language (EKI), "EKI värske uuring: paljusid tehisaru mudeleid on lihtne propagandaga mõjutada" (4 June 2026; https://eki.ee/uudis/eki-varske-uuring-paljusid-tehisaru-mudeleid-on-lihtne-propagandaga-mojutada/), and its public leaderboard (https://moodupuu.eki.ee/).

Events and current developments.

Anthropic, statement on redeploying Fable 5 (covering 30 June and 1 July 2026); "Anthropic Cyber Mission" (8 October 2026); Life Sciences Verification Program page (17 September 2026).

Ministry of Commerce and Ministry of Science and Technology of China, draft export-control catalogue and Announcement No. 48 (16 October 2026; https://www.mofcom.gov.cn/zcfb/dwmygl/art/2026/art_61446ab4beee4f63973d-4d0ab3097648.html).

Hugging Face, "Security incident disclosure — July 2026" (16 July 2026).

TechCrunch (19 October 2026) and Axios (21 October 2026) on OpenCode.

Andrej Karpathy, post on X (12 October 2026).

Arvind Narayanan and Sayash Kapoor, "A big-tent or small-tent AI safety movement?", AI as Normal Technology (1 October 2026).

Taipei Times on Taiwan's private AI computing centres (19 June 2026).

Bill Gates, "The turbulent AI era is here. The choices we make now are critical.", Gates Notes (26 August 2026).

OECD, Government Unstuck (2026).

Axios, "Scoop: Unvetted AI agents are accessing U.S. bank accounts" (22 October 2026).

Ministry of Commerce of China, spokesperson Q&A on Announcement No. 48 (16 October 2026).

Mark Carney, "Principled and pragmatic: Canada's path", Davos (20 January 2026).

Sinéad Bovell, "We Have Lost the Ability to Have a Rational Conversation About AI" (17 September 2026).

Anthropic, customer case study "Lovable"; The State of Brand, "Lovable Built $400M in ARR on Someone Else's Model" (23 April 2026).

AMD, "AMD Completes Acquisition of Silo AI" (12 August 2024).

Synergy Research Group, "European Cloud Providers' Local Market Share Now Holds Steady at 15%" (24 July 2025).

Google, "Gemini 4 Carbon: our next frontier model" (28 October 2026; https://blog.google/innovation-and-ai/models-and-research/gemini-models/introducing-gemini-4-carbon/).

Bloomberg Law, "OpenAI Founding Member Andrej Karpathy Takes Role at Anthropic" (May 2026), cited for his affiliation.

Ott Velsberg, LinkedIn post on the University of Tartu study of Estonia's AI Leap (5 September 2026).

David Sacks, post on X (18 October 2026; https://x.com/DavidSacks/status/2111867605388493282).

Keegan McBride, LinkedIn post on the Commission's technology sovereignty package (June 2026; https://www.linkedin.com/posts/keeganmcbride_today-the-european-commission-released-its-activity-7468010978568376321-WyFA).

European Data Protection Board and European Data Protection Supervisor, Joint Opinion 2/2026 on the Digital Omnibus (10 February 2026).

Andrius Sytas, "Latvia PM says Baltic Sea optical cable has been damaged", Reuters (4 January 2026), republished by ThePrint (5 January 2026).

Mistral AI, "Mistral Large 4 (Le Chonk) is now open weight" (October 2026).

AFP, "Canada buying stake in AI firm Cohere to bolster tech sovereignty" (29 October 2026).

Prime Minister of Canada, "Canada announces landmark investment in Cohere to scale secure sovereign AI" (29 October 2026; https://www.pm.gc.ca/en/news/news-releases/2026/10/29/canada-announces-landmark-investment-cohere-scale-secure-sovereign-ai).

Adam Cohen, "What It Takes to Lead in AI", OpenAI Global Affairs (29 October 2026; https://openaiglobalaffairs.substack.com/p/what-it-takes-to-lead-in-ai).

EU law and infrastructure.

European Commission, proposal for a Cloud and AI Development Act, COM(2026) 502 (3 June 2026); "Data Act explained".

EU Action Plan on Cable Security (21 February 2025); Bruce Jones, Bulletin of the Atomic Scientists (13 February 2026); Jane Ruffino, Internet Society Pulse (22 November 2024).

Regulation (EU) 2024/1689 (AI Act); Regulation (EU) 2023/2854 (Data Act); Regulation (EU) 2016/679 (GDPR).

Directive 2014/24/EU on public procurement, Article 67 (https://www.legislation.gov.uk/eudr/2014/24/article/67/data.xht).

CLOUD Act, 18 U.S.C. § 2713 (https://www.law.cornell.edu/uscode/text/18/2713).

EFTA, EEA-Lex records for the AI Act and the Data Act (checked 4 November 2026).

EFTA Surveillance Authority Decision 019/26/COL.

© 2026 NEW NORDICS AI. THIS WORK IS LICENSED UNDER CC BY 4.0 (HTTPS://CREATIVECOMMONS.ORG/LICENSES/BY/4.0/).