



REPUBLIC OF ESTONIA
MINISTRY OF JUSTICE
AND DIGITAL AFFAIRS



Nortal

TAL
TECH

National Cyber Resilience in the Age of AI

A doctrine for asymmetric cyber defence

National Cyber Resilience in the Age of AI

A doctrine for asymmetric cyber defence

Choose what must not fail. Decide where trust comes from.
Raise the floor. Defend at machine speed. Mobilise the country.
Tell the truth first.

Why this paper

The economics of cyber attack have shifted against the defender. AI has made many offensive tasks cheaper, faster and easier to scale, while the defender's duty to be lawful, reliable and right has not changed. Estimates of the annual cost of cybercrime now run into the trillions. Most nations now face an adversarial volume larger than their defensive capacity. They are, in effect, the smaller force.

Estonia has been in that position since 2007, when it absorbed the first national-scale cyberattack, and has defended against sustained campaigns by a far larger adversary ever since. The doctrine that emerges is asymmetric by design: lower the cost of defence, raise the cost of attack, and deny the political payoff that makes a country worth attacking in the first place. The aim is not invulnerability but to make incidents non-decisive.

This paper sets out that doctrine, for an era in which attacks are cheaper, faster and more numerous, and human-paced defence is no longer enough.

Authors

Co-authored by an expert working group from government, the private sector and academia:

Andres Raieste, Global Head of Public Sector, Nortal

Tõnu Grünberg, Deputy Secretary General for Digital Infrastructure and Cyber Security, Ministry of Justice and Digital Affairs

Joonas Heiter, Director General, National Cyber Security Centre (NCSC-EE) and Information System Authority

Andri Rebane, Director of Information Security Department, Estonian IT Centre

Dr Taavi Viilukas, Head of National Cyber Security, Ministry of Justice and Digital Affairs

Madis Tapupere, Chief Technology Officer for Core Banking, Luminor

Toomas Vaks, Cyber Risk Manager, Swedbank

Priit Liivak, Chief Government Technology Officer, Nortal

Andres Kütt, Supervisory Board Member, Estonian Internet Foundation

Dr Rain Ottis, Professor of Cyber Operations, TalTech

Citation: Raieste, A., Grünberg, T., Heiter, J., Rebane, A., Viilukas, T., Tapupere, M., Vaks, T., Liivak, P., Kütt, A., & Ottis, R. (2026). National cyber resilience in the age of AI: A doctrine for asymmetric cyber defence.

Table of contents

Executive summary	5
The proliferation of offensive cyber capability	8
Cyber conflict on a national scale	12
Cyber resilience policies in the age of AI	15
Policy 1. Designate the Minimum Viable State	17
Policy 2. Harden the National Trust Backbone	21
Policy 3. Raise the security baseline through enforceable standards	26
Policy 4. Defend at machine speed, lawfully	29
Policy 5. Mobilise total cyber defence	32
Policy 6. Defend democratic trust through proactive transparency	37
What good looks like: A cyber-resilient nation	41
Bibliography	43

Executive summary

Frontier AI has changed the economics of cyber conflict. The window between a software flaw becoming public and someone exploiting it has collapsed from months to days; in some cases, to hours. The work of an attack (the reconnaissance, phishing, and writing of the exploit) is being done by machines at the same time, and on the same infrastructure that powers the rest of the digital economy. Defender effort rises in a straight line, making spend-for-spend strategy unsustainable. The only path is to make every unit of defence more effective.

That market is now what democracies face. A sustained campaign against a parliamentary state is rarely fought to destroy. It is fought to make a country look broken to its own citizens, and the government slower than the news cycle in answering for it. Every move is kept just deniable enough to keep the response national and slow, rather than allied and fast. Damage is the means; the erosion of public trust is the end. The campaign sits in the gap between what a free press can plausibly report and what an open court can prove, and the defence must operate inside that gap or arrive too late to matter.

With the ongoing proliferation of offensive cyber, the policy question is where national effort has the most impact on the economics of cyber, to keep the defence effort sustainable and actionable against AI-era threats. This paper sets out a doctrine of six moves in the order a state should make them.

Choose what must not fail. Designate the Minimum Viable State: the handful of functions that if lost, harm life, financial system, public order or the workings of the state, and rehearse the fallback for each. **Decide where trust comes from.** Anchor digital trust in a set of services hardened to a level past the point where breaking them is affordable. **Raise the floor.** Close the cheap path with a baseline every regulated entity must meet, a stricter regime above it for the systems the country cannot lose, and a regulator with the authority, and the will, to suspend or force fallback mode of a service that is unsafe to fly, forcing adversaries to spend more on attacks. **Defend at machine speed.** See the estate, shrink it, and give the defender the legal authority, granted in advance and reviewable after the fact, to act within the hours an attacker now has. **Mobilise the population.** Plan cyber defence the way the Nordics plan total defence: as one national workforce that runs from the professional, through the reservist and the student, to the citizen, so that capacity scales with the

threat. **Tell the truth first.** Trust survives openness, not suppression, and becoming proactive in transparency is the easiest means to deny the aggressor the political payoff.

The outcome is not invulnerability. No state, however well governed, will keep out every attack. The real outcome is attacks becoming non-decisive: where incidents stop becoming crises, crises stop becoming national paralysis, and an attacker pays more than the result is worth. A citizen, watching the system bend, still trusts that it will hold.



FireEye CEO Kevin Mandia, SolarWinds CEO Sudhakar Ramakrishna and Microsoft President Brad Smith before a Senate Intelligence Committee hearing on the SolarWinds breach, Washington, February 23rd 2021. The breach compromised nine federal agencies and roughly 18,000 customers, turning a software compromise into a systemic national-security problem.

Photo: Drew Angerer / UPI

The proliferation of offensive cyber capability

For a decade, the economics of offensive cyber drifted in one direction. Tools grew cheaper, marketplaces deeper, and freelancers more capable. Mounting a serious attack no longer required a state intelligence service; a competent operator and a credit card would do. Defenders adjusted, and the system held. Damage attributed to cybercrime now runs into the trillions of dollars annually.¹

The cost curve has shifted

Artificial intelligence has changed the curve. The drift of a decade is being compressed into a far shorter cycle, and the capability behind each price point has deepened. Reconnaissance, social engineering, vulnerability discovery and exploit development, work that once rationed serious offensive cyber to a small cadre of skilled operators, are being automated and sold as off-the-shelf services. Attacks have become cheaper, faster and available to a much wider set of actors. The shocks of the past two decades, from Estonia in 2007 to WannaCry, NotPetya, SolarWinds and the war in Ukraine, drew regulatory adjustments, but none shifted the cost curve. This has.²

Defenders entered the AI era with practices designed for a slower threat: annual penetration tests, quarterly patch cycles and four-eyes controls written for one human checking another. Both speed and cost have moved against them.

The window between a software flaw becoming public and a working exploit being in circulation has collapsed from months to days, with weaponisation now sometimes preceding disclosure.³ Around half of the zero-days observed in the wild in 2025 targeted enterprise technologies,

¹ Industry estimates, e.g. Cybersecurity Ventures, Official Cybercrime Report (2024). On cyber risk as a macrofinancial-stability concern, see International Monetary Fund, Global Financial Stability Report, April 2024, Chapter 3: "Cyber Risk: A Growing Concern for Macrofinancial Stability".

² On Estonia 2007, see Ottis, Analysis of the 2007 Cyber Attacks against Estonia (NATO CCDCOE, 2008); on WannaCry, see UK National Audit Office, Investigation: WannaCry cyber attack and the NHS (2017).

³ Mandiant, M-Trends 2019 and How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends (October 2024); Verizon DBIR 2018–2024. Disclosure-to-exploitation averaged around two months across widely exploited flaws in 2018–19; by 2023 the average had fallen to five days, with weaponisation now sometimes preceding disclosure.

with edge security and networking equipment (the perimeter on which most of the rest depends) taking the largest share.⁴ Different data series point the same direction: defenders are being asked to remediate within a window the attacker can close in an afternoon.

	2018	2026
Flaw to observed exploitation	Months or weeks	Days, sometimes hours or pre-disclosure
Marginal cost per attempt	Expert-months	Compute-hours
Attacker's play	A few targets, picked carefully	Many targets at once, cheap
Who can run it	A skilled elite unit	A broader criminal base

Table 1. How the economics of attack have shifted, 2018–2026.

The marginal cost of an attempt has fallen by orders of magnitude. General-purpose AI can now perform substantial portions of the intrusion workflow, from the reconnaissance that opens a campaign to the localisation that adapts it for the target market. The decisive effect is parallelism: many tailored attempts running at once, against many targets, on the same infrastructure that powers the rest of the digital economy.⁵ The first campaigns combining several intrusion steps under AI direction were disclosed in late 2025; the structural shift holds whatever line is drawn on machine autonomy.⁶

TTE (Time-to-Exploit) measures the gap between CVE disclosure and confirmed exploitation




Figure 1. Collapse of the patch window, 2018–2025 (order-of-magnitude trend). Source: Mandiant and Google Threat Intelligence Group time-to-exploit data, visualised at zerodayclock.com.

⁴ Google Threat Intelligence Group, 2025 zero-day exploitation review (2026).

⁵ UK AI Security Institute and frontier-lab evaluations (Anthropic capability evaluations under its Responsible Scaling Policy, OpenAI red-team reports, Google DeepMind frontier-safety evaluations) report substantial reductions in the marginal cost of intrusion sub-tasks under AI-assisted workflows.

⁶ Anthropic, Disrupting the first reported AI-orchestrated cyber espionage campaign (November 2025). Independent commentary contested the degree of model autonomy in the specific campaign.

The asymmetry runs deeper than tooling

The defender's costs do not move on the same curve. An attacker only needs one path to succeed and can tolerate a high error rate, since a wasted attempt costs almost nothing. A defender must cover every path and cannot tolerate the same error rate, because a false positive isolates a hospital or freezes a payment rail. Many consequential defender actions, such as blocking, attributing, sanctioning or taking a service offline, are acts of public authority. The requirement that they be lawful, reversible and explainable is where defender AI hits its ceiling.

The asymmetry is also uneven across the digital estate. It hits the systems hardest that are easy to attack and slow to patch, such as industrial controllers, medical devices or the legacy estate that runs much of government and banking. Its bite is weaker on systems that can be tested, fixed and re-deployed at machine speed, such as browsers, mobile platforms and modern cloud-native services. In between sit complex distributed systems. Their flaws are easy to find in code but hard to verify in practice. A strategy that treats the whole problem as one will misallocate effort. The chapters that follow distinguish what the state can re-engineer from what it can only contain and set policies accordingly.

Proliferation and systemic risk

Offensive cyber capability now sits in a single market that runs from state intelligence services at one end to individual freelancers wielding AI tools at the other, with contractor companies, ransomware affiliates and information-operation shops in the middle. Their boundaries are deliberately blurred, and capability has deepened in each tier. The legal categories that govern state responsibility were drafted for a physical world in which the aggressor's hand was easier to see. The result is a permissive market for sustained pressure that does not cross the threshold for the use of force.

The shift also turns cyber risk from a technology problem into a systemic one. The same tooling, applied in parallel, can target the common software, cloud platforms, payment rails and managed-service vendors on which many essential entities depend. The International Monetary Fund has begun to treat AI-era cyber risk as a question of financial stability; the European Central Bank has warned that AI can chain individually minor vulnerabilities into serious attacks faster than the conventional patch cycle can respond. The campaigns that this market now supports, and the political effects they may produce, are the subject of the next chapter.⁷

⁷ International Monetary Fund, Financial stability risks mount as artificial intelligence fuels cyberattacks (2026); see also European Central Bank, F. Elderson on AI and operational resilience in banking (2026).



The war in Ukraine showed how cyber operations, infrastructure attacks and information operations now run alongside conventional military force.

Photo: Josef Cole / U.S. Cyber Command

Cyber conflict on a national scale

The last chapter described economics. This chapter describes the purpose of that market in this economy. A sustained cyber campaign against a parliamentary democracy is rarely fought with the intent to destroy. It is fought to make a country look broken to its own citizens, and to make the government respond slower than the news cycle. Every piece is kept just deniable or ambiguous enough to ensure the response stays national and slow rather than allied and fast.

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

Four effects, in order

The campaign pursues four effects; in roughly the order it can deliver them. Services are disrupted until citizens feel the inconvenience. The information environment is distorted until the press office responds rather than announces. Officials are demoralised until decisions slow and statements weaken. These three effects lead to the fourth: the erosion of trust. A parliamentary democracy holds its legitimacy on a single condition: that citizens believe the machinery of the state, its institutions and digital systems, works as advertised. Sustained pressure attacks both beliefs. Doubt accumulates among enough citizens, journalists, allies and markets to alter what the country's government seems able to do. The campaigns fall just inside the gap between what the free press can plausibly report and what an open court can adjudicate.

Two methods: impersonation and cascade

Attackers rely on two operational methods to achieve these four effects. The first is impersonation: the aggressor forges the speech, signatures and records by which a government is believed. The second is cascade: the aggressor chooses targets in such a way that a single break propagates through every system that depends on it. Impersonation manufactures the doubt; cascade scales it. The cheapest path to national-scale effect is therefore a slow accumulation of disruption, distortion and impersonation, sequenced to compound and erode public trust, rather than a single spectacular intrusion.

Work is divided across the tiers introduced earlier. The aggressor state buys depth from its own offensive units, breadth from contractor companies that front as commercial firms, and volume and noise from the criminal and hacktivist market that sits below them. Each tier hands work to the next, and each adds a layer of deniability the next tier inherits.

Table 2 illustrates seven instruments that recur across the categories of campaigns seen in recent years, none of them new in kind. In any specific incident, attribution is probabilistic and contested.

Instrument	Most likely operator	Objective
1 Supply-chain compromise. Takeovers of open-source libraries deep in build pipelines; breaches at managed-service vendors and identity providers; software updates poisoned at the publisher (SolarWinds is the canonical public template; the XZ Utils backdoor of 2024 is widely assessed to be the work of a patient state-aligned actor).	State offensive cyber units; occasionally contractor fronts	Pre-position for cascade through one shared component
2 Industrialised exploitation. Increasingly AI-assisted vulnerability discovery against legacy estates; AI-tuned phishing and credential stuffing at machine speed; ransomware deployed through pre-staged access traded between brokers.	Contractor fronts; criminal affiliates; freelancers	Industrialise initial access; harvest the long tail of weak baselines
3 Disruption of designated services. Ransomware against hospital trusts, municipal services and other designated essential providers.	Ransomware-as-a-service crews; access brokers; state-tolerated affiliates	Disrupt; force public attention; provide cover for state operations
4 Denial-of-service and data wipers. Distributed denial-of-service against payment rails and data wipers timed to a political deadline.	Hacktivist fronts and criminal affiliates for denial-of-service; state units for wipers timed to a deadline	Disrupt around political deadlines; signal capability
5 Identity and impersonation. High-fidelity synthetic audio and video of senior officials; forged official communications and spoofed letterheads; AI-generated personas with months of plausible commit history; phishing tuned to local dialects.	State units for high-value targets; freelancers for volume	Distort; impersonate consequential decisions
6 Information operations. Coordinated inauthentic networks; fabricated leaks dressed as hacktivist disclosures; deepfake clips laundered through tabloids.	State influence units, typically via contractor shops and hacktivist personas as cut-outs	Distort; demoralise; manufacture domestic controversy
7 Hybrid pairings. Cable sabotage, energy disruption, manipulated migration and economic coercion paired with cyber, forcing the target to decide whether several incidents are one or multiple acts.	State intelligence and military units; proxy or allied states; civilian-flagged assets used as cover	Compound ambiguity; keep the response national rather than allied

Table 2. The instruments used in a sustained campaign, their likely operators, and key objectives.

AI is a multiplier for each instrument. The first reported case of an AI system orchestrating much of the intrusion lifecycle with limited human direction was disclosed in late 2025; whether that case is considered the inflection point, the cost shifts that would make such a campaign profitable have already happened.

Why the law arrives late

The whole apparatus is calibrated to fall just short of any treaty article that would justify collective response, including the use of armed force. Article 51 of the United Nations Charter recognises the right of (collective) self-defence against an armed attack. NATO has recognised that the impact of significant malicious cumulative cyber activities might, in certain circumstances, be considered as amounting to an armed attack and could lead to the invocation of Article 5 (the collective defence clause based on UN Charter Article 51) of the North Atlantic Treaty on a case-by-case basis. In practice, the threshold is high and case-specific, and sub-threshold campaigns are engineered precisely to stay below it: each incident ambiguous on its own, and impossible to add up in time to act. Operations are routed through proxies whose contractual relationship to the state can be acknowledged, denied or redrawn as the diplomatic season demands. Infrastructure sits in jurisdictions outside the relevant mutual legal-assistance treaties. Indictments name individuals who will not travel, and sanctions hit shell companies. The legal instruments work slowly, and on the wrong actors.⁸

Coordination is what makes response instruments timely. Response instruments are slow individually and only bite when several states act together: joint attribution, coordinated takedowns of state-proxy and ransomware infrastructure, sanctions against the financial rails that pay them, and supplier discipline on offensive AI capability.

The campaign is engineered to sit between what the free press can realistically report and what a court can adjudicate. The defence must operate in that gap, or it will arrive too late and be too narrow to matter. The doctrine that follows is built for that gap.

⁸ NATO, Brussels Summit Communiqué, paragraph 32 (14 June 2021). On the underlying body of doctrine, see also M. N. Schmitt (ed.), *Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations* (Cambridge University Press, 2017).

Cyber resilience policies in the age of AI

National cyber resilience is a state's ability to maintain essential democratic, economic and public-service functions during and after cyber disruption, while preserving legal accountability, public trust and coalition interoperability.

Chapter 1 set out how frontier AI has changed the cost curve of offensive cyber work, which has led to the proliferation of offensive cyber capabilities. Chapter 2 set out the political objectives such campaigns are designed to achieve; rarely territory or money alone. The prize is the erosion of public trust in the institutions on which a free society depends, and most of the apparatus is calibrated to win that prize while staying below the threshold at which collective response would be lawful.

The policy question, then, is not how to defend everything against everyone. It is where national effort has the greatest effect on the economics of attack and defence, and where it can deny the attacker the political prize even when a technical attack succeeds.

That framing is familiar. States have long defended themselves against large adversarial forces by choosing their ground rather than matching mass. They decide in advance what must not fall, prepare it deliberately, train the population for its role, contest the information space, and bind themselves into alliances that make isolated coercion expensive. The same logic now applies in the digital domain, where the asymmetry between a well-resourced attacker and an ordinary defender has widened, and the volume of plausible attackers has risen at the same time.⁹

The six policies that follow apply that logic to the digital estate. They choose the ground (Policy 1, the Minimum Viable State), harden the points through which trust is exercised (Policy 2, the National Trust Backbone),

⁹ For a longer treatment of this argument, see A. Raieste, A. Rebane, M. Tapupere and K. McBride, *Government Resilience in the Digital Age* (Oxford Internet Institute, Estonian IT Centre and Nortal, Tallinn, 2024).

raise the floor under everything else through enforceable standards (Policy 3), defend the resulting estate at machine speed and within the law (Policy 4), mobilise the country as a single defensive system (Policy 5), and protect the political prize directly by telling the truth first (Policy 6). They are intended to be read as one doctrine rather than six initiatives, each addressing a part of the same problem, and the gaps between them are where adversaries will look first. Table 3 summarises the effect of the policies.

Policy	What it covers
1. Designate the Minimum Viable State	<i>Choose what must not fail.</i> Name the functions that if lost harm life, financial system, public order or the state, and rehearse a fallback for each.
2. Harden the National Trust Backbone	<i>Decide where trust comes from.</i> Anchor digital trust in a set of services hardened to a level past the point where breaking them is affordable.
3. Raise the security baseline through enforceable standards	<i>Make security an operating licence, not paperwork.</i> Close the cheap path with a baseline for all, forcing adversaries to spend more on attacks.
4. Defend at machine speed, lawfully	<i>Give defenders the tools, and the law, to keep up.</i> See the estate, shrink it, and let the defender act inside the attacker's window rather than after it.
5. Mobilise total cyber defence	<i>Train the country, not only the cyber team.</i> Spreads defence across a national workforce, so that capacity scales with the threat.
6. Defend democratic trust through proactive transparency	<i>Tell the truth first.</i> Trust survives openness, not suppression, and becoming proactive in transparency is the easiest means to deny the aggressor the political payoff.
<i>Lower the cost of defence · Raise the cost of attack · Deny the political payoff</i>	

Table 3. The policy architecture. Together, these make attacks non-decisive: a country that fails well, where an attacker pays more than the result is worth.

POLICY 1.

Designate the Minimum Viable State

Choose what must not fail

A state that defends every asset to the same standard ends up defending none of them well. This fails in two ways. The first is cascade: an attacker finds the weakest peripheral system and rides it through shared identities and dependencies in the systems that the country cannot lose. The second is sovereignty: a capability the state does not own can be withdrawn, throttled or repriced by a foreign supplier, often on the day the state needs it most. Both failures come from the same missing decision by the state: which functions must remain operational under sustained attack or supply-chain failure, and which may be allowed to degrade.

Tier the estate

It helps to name three tiers (Table 4). At the top sits the Minimum Viable State: the functions the country cannot lose for an hour. Below it are essential services that can tolerate degradation for hours or days. Below that, the wider economy, held to a national hygiene floor.

A function belongs in the Minimum Viable State only if its loss within minutes or hours would directly threaten life, constitutional continuity, public order, sovereign authority, financial settlement, or national-scale trust. A politically important service is not automatically MVS. It becomes MVS only when there is no acceptable degraded mode.

The Minimum Viable State belongs in law, or at least in a national strategy. The categories can be public; technical detail need not be. Once a function sits in the top tier, the standard, procurement discipline, monitoring and rehearsed fallback that go with it follow automatically; nothing stays in the top tier on a self-assessed exemption.

Tier	What it covers	Examples	Tolerance	Standard
1. Minimum Viable State	Functions which if lost, harm life, the financial system, public order or the constitutional process	Population registers and national identities; court and property registers; key telco systems, electric-grid and water control; hospital and emergency-line systems; central-bank settlement and core payment rails	Minutes	Continuous independent audit; rehearsed fallback; statutory inclusion
2. Regulated essential services	Sectors the country depends on but can run degraded for a working day	Wider health estate; taxes and benefits; customs; transport and ports; telecom operators; large banks beyond settlement; energy and water firms	Hours to days	NIS2-style obligations; statutory audits; attack-surface budget; published recovery targets
3. Wider economy	The long tail of which individual failure is tolerable but a collective compromise feeds the cascade	SMEs; local-government back offices; consumer digital services	Best effort, with limits	National hygiene floor enforced through procurement, insurance and consumer-protection laws

Table 4. Different tiers of the state and the wider economy

Source what matters

For every top-tier function, the state needs to know who is on the hook on the worst day. There is a small menu of answers: run it in-house, regulate a domestic supplier, pool the supply with allies under treaty, or buy it commercially with audit and exit rights written in. Each option carries its own risks, sovereignty choices, and recovery story. The same logic applies to frontier AI: in this context, the strategic asset is not the model itself but the bespoke systems and integrations around it. What no state can afford is a top-tier function for which answers have never been chosen, only became the historical default. Behind the law sits a plain register: for each top-tier function, what it depends on, who supplies it, where the supplier sits, and what the country falls back to if the supply breaks. Where an essential service rests on something the state does not run itself, the state keeps the right to inspect and the duty to explain.

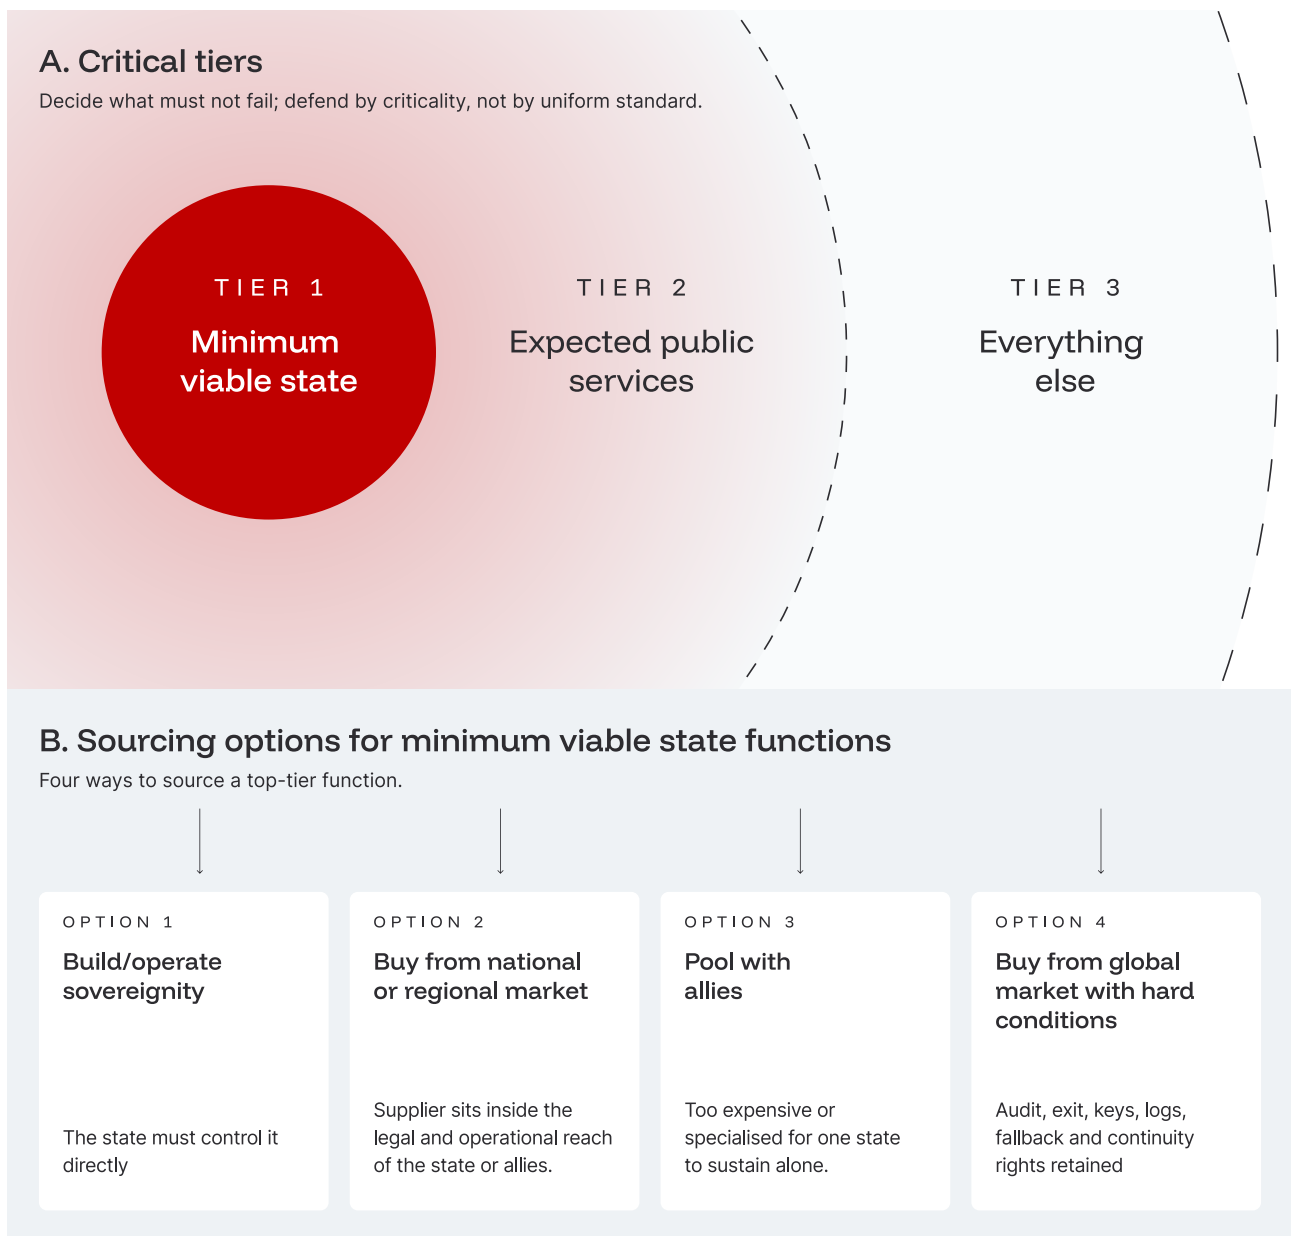


Figure 2. Different sourcing modes in the minimum viable state.

Rehearse the fallback

Untested fallback is not a fallback; it is a hope. Every top-tier function needs an alternative the state has run, whether on paper, in parallel or pooled with allies, and on a schedule the public can see. Rehearsal is what turns a line in a register into something the country can reach for on the day a primary supplier disappears. The results go to the regulator, and when they matter, to the public.

Key recommendations

Ask: Which national functions must remain operational under sustained attack, is that list realistic and capped, and is that choice fixed in law?

-
- | | |
|-----|--|
| 1.1 | Designate the Minimum Viable State in statute. Cap its scope to a realistic level. |
|-----|--|
-
- | | |
|-----|---|
| 1.2 | Settle the sourcing mode for every top-tier function strategically, considering fall-backs, redundancy, commercial availability and sovereign control as key decision points. |
|-----|---|
-
- | | |
|-----|---|
| 1.3 | Keep a live register of what each top-tier function depends on and what the country falls back to if the supply breaks. |
|-----|---|
-
- | | |
|-----|--|
| 1.4 | Treat any fallback that has not actually been run as a hope and rehearse each one on a published schedule. |
|-----|--|
-

POLICY 2.

Harden the National Trust Backbone

Decide where trust comes from

In a digital state, trust is national infrastructure, not a property each system has to invent for itself. Whether a ministerial notice is genuine, whether a contract was really signed by the person named, whether a property register entry is the one the law recognises: these questions should have the same answer, drawn from the same hardened foundations, whether the party asking is a citizen, a bank, a regulator or a piece of software acting on behalf of any of them. The country needs a federated family of trusted services (identity, signature, registers, exchange, audit, hosting) through which every consequential act inherits its authority. Call it the National Trust Backbone. The Backbone is defined by assurance rather than provenance: it is the set of services, public or private, that the country cannot afford to lose trust in, held to one top-tier standard. A national identity service (or other elements of digital public infrastructure), an interbank clearing rail and a critical legal registry may all sit inside it; what they share is the standard, not the sector. Name it in law, fund it, federate it, and let the rest of the state, and the wider economy, lean on it. Where Policy 1 chose the functions that must not fail, the Backbone is the family of services, owned by various parties under a common regulation, from which those functions inherit their legal and digital authority.

Inherit trust from hardened services

The argument is that digital trust should be inherited from hardened national foundations, rather than reinvented system by system. The backbone is a legally recognised family of federated trust services, held to one top-tier assurance standard, with redundancies and fallbacks. A set of components, audited to the top standard, reaches a level of assurance no patchwork of sectoral half-builds will match. Some solutions are commercial, state-owned, or pooled with allies; the model varies, the standard does not. Table 5 shows the key layers in the trust backbone.

The same logic makes the backbone the country's highest-payoff target: the more it carries, the more an attacker gains from breaking it. The answer is not to scatter trust thinner, but to spend more on the federated set already chosen: segmentation, cryptographic agility and post-quantum migration, independent audit and rehearsed fallback. Concentrated investment in plural, hardened services makes the attacker's bet unprofitable.

Layer	Function
Digital identity	Resolves who is acting and the entitlements they hold: citizen, official, company, system or AI agent.
Qualified signature and authorisation	Non-repudiable evidence that an authorised action took place, verifiable by any recipient in seconds.
Authoritative registers	Hold the records on which decisions rely, with cryptographic integrity over every consequential change; signed at source.
Secure exchange, interoperability	A shared channel for secure, federated data exchange or transactions, replacing bilateral integrations and pooled copies.
Audit, transparency and redress	Citizen-accessible audit log of who queried personal data, parliamentary oversight, independent regulatory access, and route for citizens to challenge misuse.
Protected hosting	Regulated cloud capacity for essential workloads, under common identity, logging and key management.

Table 5. The key six layers of the National Trust Backbone.

Containing the blast radius: graceful backbone failure

The backbone must be designed on the assumption that a component will eventually be compromised, and that the damage must be contained at the layer where it occurs. Estonia's 2017 ID card chip flaw, which forced the precautionary suspension of certificates on some 750,000 cards, showed why: the country kept functioning because it had more than one way to prove who you were. Four principles make this possible.

Redundancy where it matters most. No critical trust function should depend on a single product or cryptographic family. A mature digital state runs several means of authentication in parallel issued under one legal regime but built on different stacks. The same logic applies to signature, time stamping and the hosting of consequential workloads.

Fallbacks that have actually been used. A second route that has never carried real traffic is a hope, not a fallback. The backbone's alternatives should be in routine use and rehearsed on a published timetable; switchover targets are measured in days.

Integrity that travels with the record. Authoritative registers are the country's memory of legal facts. Every consequential entry is signed at source, chained to its predecessors, and verifiable by any recipient without trusting the operator. A successful intrusion then becomes an event that can be detected, bounded and reversed, rather than one that has to be believed or denied.

Separation that prevents propagation. Registers are kept apart, each holding only the records its regulated function requires. The service that signs is separated from the service that issues identities. Data moves through a controlled exchange channel rather than by handing out copies, and the most important registers have a second home, across redundant clouds, or in a friendly country under a treaty. Backbone identity and signature are the ordinary route for any consequential act, including those taken by software agents, whose credentials are issued for the job and the moment and revoked when those are done. Standing privileges age into liabilities.

A backbone built this way is still not invulnerable, but it can absorb a compromise and keep the state running — which, at national scale, is the only useful definition of hardened.

Key recommendations

Ask: Can a citizen verify, in real time, whether a public action, record, instruction or signature is authentic, and can they still do so even if part of the backbone itself is compromised?

-
- | | |
|-----|--|
| 2.1 | Name the National Trust Backbone in statute and fund and operate it as critical infrastructure, not a digital programme. |
|-----|--|
-
- | | |
|-----|--|
| 2.2 | Require that no single compromise in the trust backbone can take down more than one of its layers. |
|-----|--|
-
- | | |
|-----|--|
| 2.3 | Make backbone identity and signature the default route for every consequential public or regulated act, by humans and software agents alike. |
|-----|--|
-
- | | |
|-----|--|
| 2.4 | Make the trust backbone replaceable, cryptographically, before the country must replace it in a hurry. |
|-----|--|
-

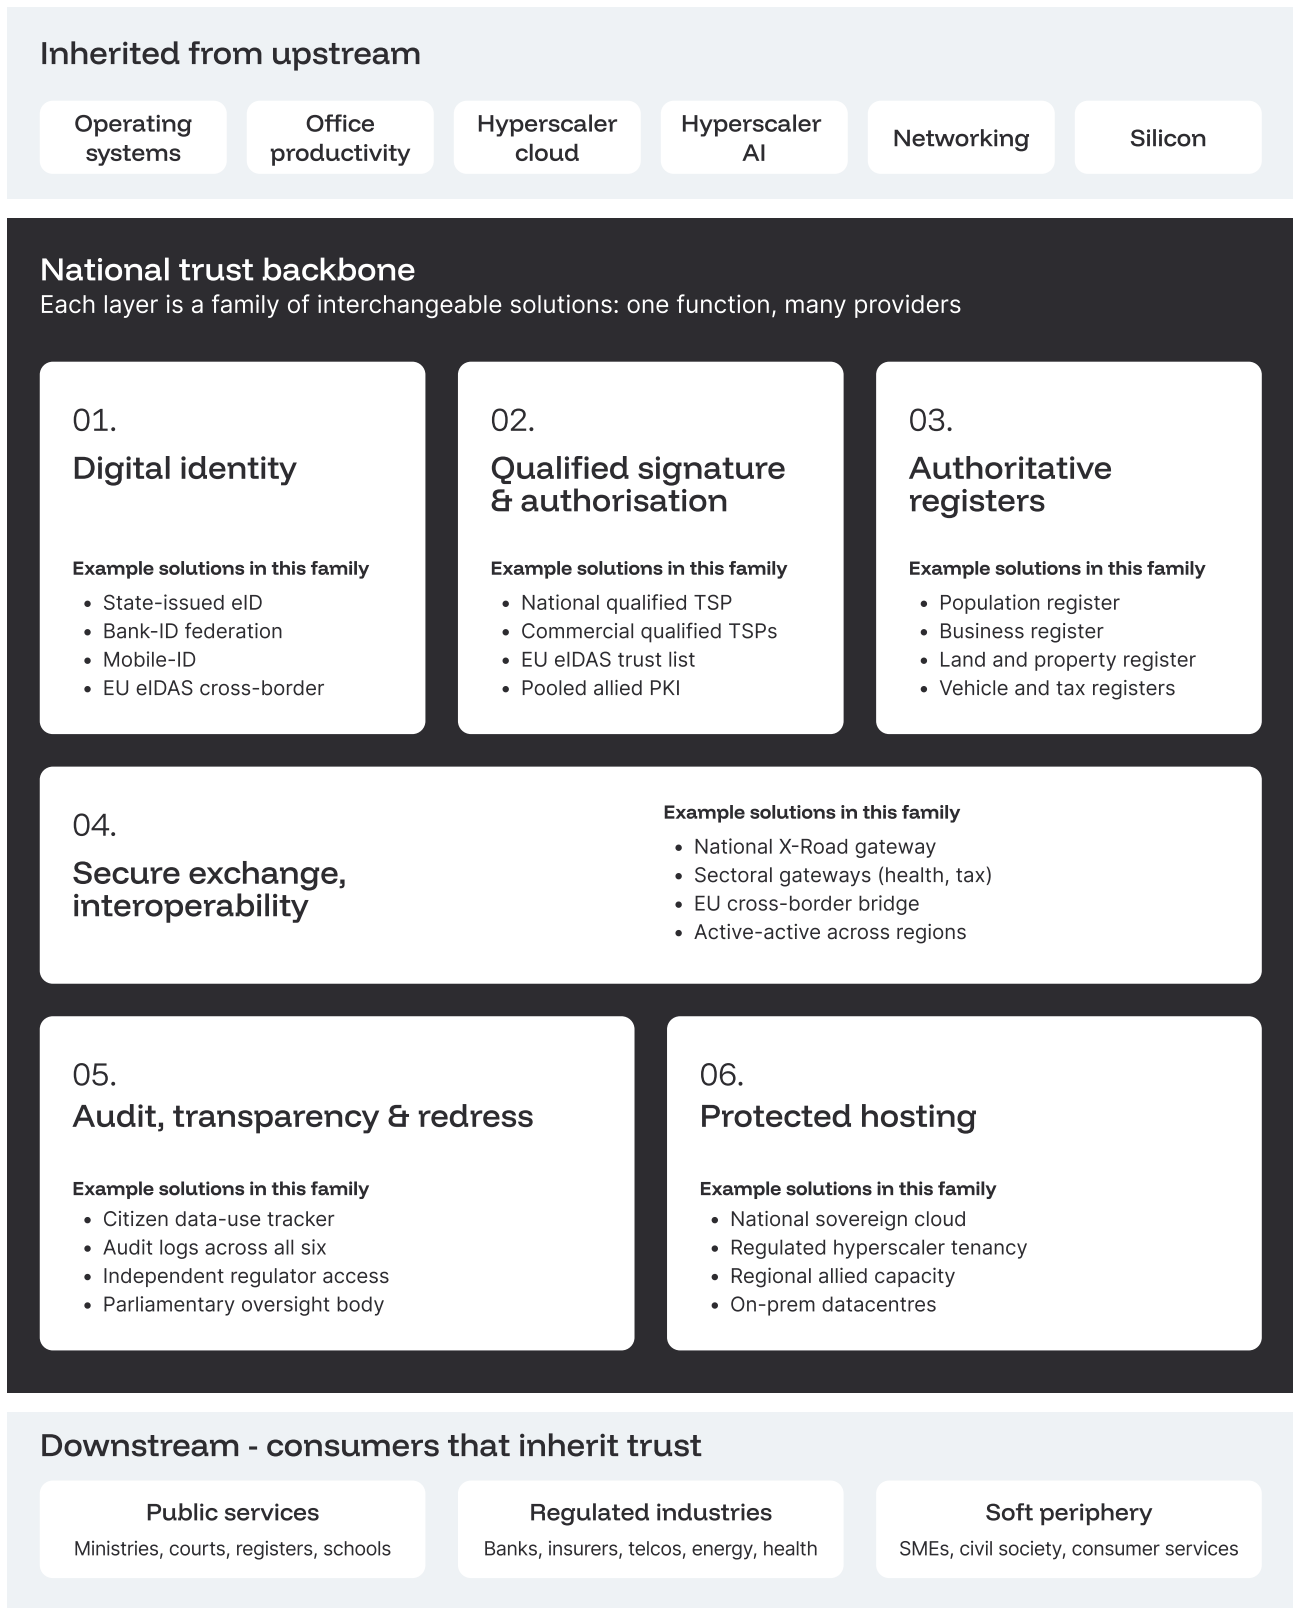


Figure 3. The backbone is a family of capabilities, inheriting capabilities from (global) upstream and passing trust downstream. Examples of what may be considered part of it differ country-to-country.

POLICY 3.**Raise the security baseline through enforceable standards***Make security an operating licence, not paperwork*

Even a patient adversary begins on a cheap foothold. Unpatched software, default credentials, unsegmented networks, copy-pasted code with known vulnerabilities: the same shortcuts that subsidise routine crime are the entry the sophisticated operator reuses before any tradecraft is spent. The attacker, criminal or state, reaches for the cheapest path that still works, because the campaign budget is finite and the next target is one scan away. Frontier-AI tooling has sharpened that logic: reconnaissance and exploit-tuning that once required a specialist now run as a workflow, in parallel, against many targets at once. The aim of a standards regime is twofold: to make the routine campaign unrewarding, and to force the patient adversary to spend tradecraft it would rather hoard, denying the cheap foothold that funds the rest of any operation.

Set a collective baseline and a stricter national regime

Risk-based standards alone no longer suffice. A short list of controls should be non-negotiable for every regulated system; a second list is non-negotiable for every new or materially updated one. The point is not to enshrine them in doctrine but to remove the discretion to skip them: where controls are missing, the system is not compliant, irrespective of how the rest of the risk assessment reads.

Most modern jurisdictions are converging on the same layered pattern: a horizontal floor for essential services, a stricter operational regime for finance, product-security obligations for what is sold to the public, and a data-protection regime underneath. Europe's regulatory stack is a good example of this, but the pattern is universal. For the very few systems the country cannot lose, the collective floor is not enough; a stricter national regime sits above it, updated faster, so an operator answers to one set of rules rather than two.¹⁰

¹⁰ Europe's layered stack comprises the NIS2 Directive (EU) 2022/2555 as the horizontal floor; DORA, Regulation (EU) 2022/2554, as the financial-sector regime; the Cyber Resilience Act, Regulation (EU) 2024/2847, as product-security law; the AI Act, Regulation (EU) 2024/1689; and the Cyber Solidarity Act, Regulation (EU) 2025/38. The European Digital Identity Framework, Regulation (EU) 2024/1183 amending eIDAS, sits alongside as the identity layer.

Enforce visibly

A standard becomes serious through audits and consequences. Self-attestation is not an audit. Statutory independent audits at a cadence that tracks criticality feed a regulator with the staffing, authority and willingness to impose costs on non-compliant organisations. A regulator that does not impose costs is itself part of the attacker's subsidy.

The aviation regulator's instrument applies here: the authority can ground the fleet. The cyber equivalent is the power to ground a system. For the few services the country cannot lose, serious non-compliance must trigger suspension or a forced switch to a tested fallback: from one identity provider to another, one settlement rail to another. The power must be in statute, time-limited, and reviewable in court. The knowledge that the regulator can ground a system changes the price of non-compliance, not how often it is used.

Give procurement bite

A procurement clause with bite costs less than a regulator that must step in repeatedly. What the state buys should default to the tougher standard: hardened identity, signing where records are made, a tested fallback, audit access written into the contract. Non-functional requirements on cyber security make it into vendor contracts. Exceptions are recorded, on the record. The state is among the largest customers in any digital economy; what it asks for sets the price of compliance for every other customer of the same supplier.

Reward balances coercion. Firms that commit beyond the baseline should earn additional advantages: eligibility for state contracts that demand hardened security, and a public certification that customers and counterparties recognise. Together, a coercive baseline, visible enforcement and incentives make meeting the standard the firm's cheapest path.

Key recommendations

Ask: When a critical digital service is found non-compliant, can the regulator force it into safe operation before an attacker uses the gap?

-
- | | |
|-----|---|
| 3.1 | Make a shortlist of non-negotiable cyber controls for every regulated system, and a stricter list of non-negotiables for the systems the country cannot lose. |
|-----|---|
-
- | | |
|-----|--|
| 3.2 | End self-attestation for systems the country cannot lose and put them under statutory independent audit. |
|-----|--|
-
- | | |
|-----|---|
| 3.3 | Give the cyber regulator the power to ground a backbone system the way an aviation regulator grounds a fleet: legislate a suspension or forced-fallback order for top-tier systems, subject to quick judicial review. |
|-----|---|
-
- | | |
|-----|---|
| 3.4 | Use what the state buys to set the price of compliance for everyone else: default every central-government framework contract to the hardened standard. |
|-----|---|
-

POLICY 4.

Defend at machine speed, lawfully

Give defenders the tools, and the law, to keep up

The defender has hours when they used to have weeks. The attackers are using AI for accelerating each step. Defence must follow suit. For this, the defenders need regulatory power to match the attack speed, see what they are protecting, shrink the attack surface, and continuously defend what remains. Each step is straightforward in principle and uncomfortable in practice; together, they are what allows the country to act before the attacker has finished.

Pre-delegated authority to match the attack speed

To match the machine speed, the defender needs the power to act — granted in advance, written into law, and reviewable in court after the fact — against systems inside the country's own jurisdiction. Action across borders is what coalitions are for. Data-protection law applies in full, and the powers must be rehearsed in joint exercises so they are not invented during the first real crisis. Where the same containment can be done through a contract clause, it should be, on the same terms.

See the estate

It is difficult for the state to effectively shrink an attack surface without first knowing what it is. The country needs a live, machine-readable map of its digital estate, including what each system does, who owns it, how it is supplied, and where it is exposed. A map drawn by humans is out of date the day after it is finished. The new work is to automatically keep the map current to actual code and traffic. The map itself is then a top-tier service, hardened like the others.

Shrink the estate

The cheapest way into a national cascade runs through what the state has stopped looking at: the old payroll system, the portal nobody documented, the open-source library three hops from anything anyone owns, the managed service whose client list reads like a directory of essential entities. Absence of an owner is grounds for retirement, not a problem to be parked for the next administration. What is kept is modernised on a published timetable, with the system named.

For the entities a regulator already watches, it should publish a budget (a cap on the doors and privileged accounts each entity is allowed to keep open) and tighten it year on year. The cap, the instrument within the perimeter an attacker reaches for, shrinks on a timetable everyone can plan for, rather than after the incident.

Defend continuously, lawfully

The old patch cycle no longer fits the threat. AI now strings small flaws together into serious attacks in an afternoon, so priority must follow what is being used against the country, not the abstract severity score. The defender's reply should be symmetric in tooling and asymmetric in authority: AI triages alerts against the estate map, correlates weak signals across sectors, scores which disclosed flaws are likely to be weaponised against the country, and, within pre-set rules, contains low-consequence incidents without waiting for a human in the loop. The state needs an authoritative list of the vulnerabilities being exploited in the wild with remediation duties at a pace the regulator can enforce. The legacy estate that carries the most risk is modernised on a published timetable, and the modernisation itself (reading undocumented code, recovering lost design intent, generating and testing fixes) is work that AI now shortens from years to months.

The AI tools used for this must also be bought with the same discipline as the systems they protect before they are trusted with anything consequential. Sanctioned tools must displace shadow AI inside the public sector: an unsanctioned tool inside the defender's perimeter is a supplier the state has not chosen.¹¹

¹¹ UK National Cyber Security Centre, guidance for defenders on frontier AI (including Why cyber defenders need to be ready for frontier AI; Preparing for a vulnerability patch wave; 10 questions to ask when using AI models to find vulnerabilities). The evaluation discipline draws on NIST AI 100-2 E2025, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations (2025).

Key recommendations

Ask: Within the attacker's window, can the state name which top-tier systems are exposed to a live vulnerability and act on the answer in law?

-
- | | |
|-----|--|
| 4.1 | Give the defender a legal right to act at machine speed inside national jurisdiction and under pre-defined, specific conditions. |
|-----|--|
-
- | | |
|-----|--|
| 4.2 | Run a live map of the national digital estate and protect it like one of the systems it watches. |
|-----|--|
-
- | | |
|-----|--|
| 4.3 | Fund and run continuous vulnerability operations and legacy modernisation, accelerated by AI, against the national digital estate map, the Minimum Viable State, or at the very least, the trust backbone. |
|-----|--|
-
- | | |
|-----|--|
| 4.4 | Cap the number of internet-facing doors and standing privileges each regulated entity is allowed and tighten the cap year on year. |
|-----|--|
-

POLICY 5.

Mobilise total cyber defence*Train the country, not only the cyber team*

Cyber resilience, like any form of national defence, must be built against the threat it faces. The principle is old: a state sizes its forces, its reserves and its civil preparations to the adversary. In most countries, cyber defence has yet to be brought fully under this discipline. It should be planned whole-of-society, with manpower, reserve capacity and citizen literacy calibrated to a measured threat curve, reported publicly and held to the same standard of accountability as the armed forces themselves. A government that cannot say whether its cyber-capable workforce is growing faster than the danger it faces cannot claim to be defending the country in any serious sense. A digital state needs cyber capacity on the scale of a national workforce, planned the way the Nordics plan total defence, not as a personnel programme bolted onto a security strategy.

Capacity layer	Meaning
Citizen cyber literacy	Population-level ability to recognise phishing, deepfakes, fake channels, unsigned documents and verification pages.
Higher-education pipeline	Vocational and university programmes producing graduates with the skills the national cybersecurity community needs; cybersecurity content embedded into adjacent disciplines (engineering, law, and public administration) rather than confined to dedicated degrees.
Professional cadre	National CSIRT, government SOC, regulators, incident responders, forensic teams, secure architects.
Cyber reserve	Defence-force or civil-reserve pathways that include cyber roles and exercise alongside the regular workforce.
SME and civil-society support	Defensible baseline tools, guidance and shared services for organisations too small to defend themselves alone.
Allied pooling	Share specialist capacity and expertise among allied states.

Table 6. Six capacity layers, planned as a single national output.



Locked Shields, NATO CCDCOE's annual cyber-defence exercise. The most scarce cyber capabilities are increasingly pooled, exercised and sustained across allied states.

Photo: Arno Mikkor / The NATO Cooperative Cyber Defence Centre of Excellence

Plan capacity as one national output

The closest precedent is the doctrine of total defence: the Nordics, Singapore and Switzerland treat military capacity as a single national workforce, in which conscripts, reservists and civil protection sit inside one plan. The cyber equivalent has the same shape and the same audiences: the professional the state must keep, the reservist it can call on, and the citizen it must teach to spot the phishing email. The university pipeline and the small-business perimeter sit between them. All of it is forecast and funded as one workforce, not separate budgets competing for staff.¹²

Retention is an executive responsibility, not a personnel-policy detail. The non-salary levers matter as much as pay: clearances that move at industry pace, rotations, and sabbaticals. Attrition should be reported annually, by role, so the public can see when the workforce is being run down.

Pool scarce roles

Pooling is how the rarest skills are made available to every ministry rather than recruited separately by each. At the national level, a country can run a single response team, security operations centre and group of incident responders, so that the specialist a ministry needs at three in the morning is reachable. Pooling is also the only credible answer to the rarest skills of all (testing the new generation of AI systems, evaluating them for misuse, migrating to post-quantum cryptography) which no single country can sustain alone.¹³

The reserve sits alongside the salaried workforce. Volunteer and reservist cyber units (military or civil), give the state surge capacity it could never afford to keep on the permanent payroll, and a route into the regular workforce. Their value lies less in headcount on paper than in the exercises they run with the response team, the regulator and the essential operators, so that on the day a sustained campaign begins, the reserve is already trained, trusted, and known.

¹² On total defence as a planning doctrine, see the Swedish, Finnish, Singaporean and Swiss frameworks for whole-of-society defence, in which cyber capacity is treated alongside military and civil reserves.

¹³ EU-CyCLONe, the European Cyber Crisis Liaison Organisation Network operated under ENISA, is the standing example of allied pooling of crisis coordination at scale.

Build cybersecurity literacy through life-long learning

Education sits underneath all of it. Basic cyber hygiene belongs in school, gets refreshed in adult life, and is easy for anyone changing careers to pick up. At the professional end, certifications and university courses must keep pace with how attackers truly work, which means the people who teach them must spend time inside the agencies and operators they teach about.

Key recommendations

Ask: Does national investment in cyber defence, in money and in people, bear any honest relation to the share of national harm that cyber attacks now cause?

-
- | | |
|-----|--|
| 5.1 | Plan the country's cyber workforce as one, not with separate budgets fighting for staff. |
|-----|--|
-
- | | |
|-----|--|
| 5.2 | Add cyber to the national service and reserve, so the state can draw on skills it cannot sustain on a permanent payroll. |
|-----|--|
-
- | | |
|-----|--|
| 5.3 | Pay teachers to spend time inside the agencies they teach about, so syllabuses keep up with how attackers really work. |
|-----|--|
-
- | | |
|-----|---|
| 5.4 | Pool the capabilities no single country can sustain alone with allies and practise them before the day of need. |
|-----|---|
-
- | | |
|-----|---|
| 5.5 | Teach citizens to check the channel before they believe the message, from school age onwards. |
|-----|---|
-

POLICY 6.

Defend democratic trust through proactive transparency


Tell the truth first

A democracy that defends trust by destroying its conditions has lost. The campaign described in the second chapter has two ways to win. It can accumulate enough public doubt to narrow what the government feels it is able to do. Or, by overreach, it can goad the defender into becoming the surveillance state the attacker was claiming it was all along. The reply has three parts: tell the truth first, rebut before the rumour is believed, and place every defensive instrument under independent oversight. People do not trust institutions because nothing ever goes wrong. They trust them when failures are reported quickly, explained clearly and investigated independently. The aviation regulator publishes a preliminary report within days and a full investigation in due course; openness, not suppression, is what sustains trust under pressure.

Disclose by default

If the state speaks first, the rumour has no legs. A legal duty by public bodies and essential operators to disclose material incidents, quickly and in terms the public can act on, turns transparency from a reputational risk into a defensive instrument. The campaign trades on the gap between an incident and the public explanation of it; a state that closes the gap inside the news cycle has already won the moment.

The templates for what the state will say are written before the incident, signed on a channel citizens know to look for, and tested in joint exercises with the regulator and the press. The same discipline applies in calmer times. A citizen can already be shown, in real time, who in the state has queried their personal data, when, and on what legal basis: a tangible piece of openness at the individual level, with a route to challenge and a parliament watching the system. Estonia has run such a log since 2017.



Strong digital identity

The same identity authenticates request and citizen.

Digital twin

Every register entry that names you, surfaced in one place.

Real-time logging

Every access leaves a temper-evident record.

Legal-basis tags

What authority, what purpose, under which law.

Consent dashboard

Opt-in controls for optional data sharing, revocable.

Escalation route

One tap to the data-protection authority, with record attached.

Enabled by the trust backbone

Identity

Registers and
integrity

Exchange and
interop

Audit and
redress

Figure 4. Illustration of transparency in the citizen's pocket: who used my data, when, under what legal basis, for what purpose, and how to push back. Based on the Estonian citizen personal-data tracker service.¹⁴

¹⁴ Information System Authority (RIA), Electronic identification and trust services (Republic of Estonia); the personal-data usage log is integrated with the state portal eesti.ee.

Pre-bunk and rebut

A standing team inside the state, working with platforms and public broadcasters, identifies and answers manipulated content before the rumour is widely believed. The trick is to explain the pattern in advance, not only the incident, so the public learns to recognise the move when it appears. Sweden has run a Psychological Defence Agency on these lines since 2022; this is the working reference.¹⁵ Taiwan complements the institutional model with an operational one: line ministries publish a signed clarification within roughly an hour of a rumour appearing, in a form short enough to travel on the same platforms, the discipline being measured in tempo and reach, not volume.

The channels on which the state speaks must be advertised and rehearsed before elections, so that a faked ministerial clip can be answered on a known channel within the same news cycle. Signed channels for ministers, election notices, and emergency instructions are advertised before they are needed, not invented during a crisis. The rules under which the state works with platforms are published, so the line between defence and censorship stays where the public can see it.

Constrain the defenders

Every defensive instrument that touches citizens should be subject to scrutiny, with the right in law to inspect it and the power to recommend its suspension. What must be prevented is the slow widening of a defensive tool into a general-purpose machine for watching the people it was built to protect. Defensive tools must be built without the surveillance features they do not need: what is not collected cannot be misused. The oversight regime is what makes the difference between a free state defending itself and a security state in everything but name. This is what makes a visible difference to the public.

¹⁵ Swedish Psychological Defence Agency (Myndigheten för psykologiskt försvar), established 1 January 2022 with a remit covering the identification of, and response to, foreign malign information influence.

Key recommendations

Ask: When a fabricated ministerial video, forged notice or service-disruption rumour appears, can citizens verify the truth through a known channel within the same news cycle?

-
- | | |
|-----|--|
| 6.1 | Set a 24-hour disclosure duty, with the template and channel rehearsed jointly with the regulator and the press. |
|-----|--|
-
- | | |
|-----|--|
| 6.2 | Designate one official channel per ministry, publish it, and rehearse it before every electoral cycle. |
|-----|--|
-
- | | |
|-----|---|
| 6.3 | Let every citizen see who in the state has looked at their data, when, and on what legal basis, in real time. |
|-----|---|
-
- | | |
|-----|---|
| 6.4 | Explain the manipulation pattern before a rumour appears, operating a standing pre-bunking and rebuttal capability. |
|-----|---|
-

WHAT GOOD LOOKS LIKE:

A cyber-resilient nation

The point of this doctrine is not invulnerability. No state, however well governed, will keep every attack out. The point is that incidents stop becoming crises, and crises stop becoming national paralysis.

A cyber-resilient nation can be recognised by how it behaves under sustained pressure. The lights stay on, payments clear, ambulances are dispatched, courts meet, elections are run and counted. Functions outside that core may degrade, but they recover on a known timetable, with the public told plainly what is down and when it will return. Officials sign what they say, and citizens have the means to check. Rumours find news cycles already full of official responses.

Behind the visible calm sits the work this paper has set out. Ministers know which functions must not fail and who is on the hook for each. The trust backbone is hardened, segmented and auditable, with consequential records signed at the source. The security floor is enforced by statute, audit and procurement, not by exhortation. The attack surface is mapped and shrunk; defender AI is bought and run with the same discipline as the systems it protects. A national cyber workforce sits inside a wider pool of reservists, students and informed citizens. Pre-delegated operational authority lets defenders act inside the attacker's window, under expedited judicial review. Allies pool telemetry, attribution and surge capacity that no single country can afford alone.

The test is whether this changes the attacker's arithmetic. A campaign that once promised political payoff at a modest cost now demands deeper capability for thinner returns. The minimum viable state continues. Supplier compromises do not cascade. Identity fraud is caught at the signing layer. Zero-days are remediated faster than they are weaponised. Disinformation is spread in a state that has already spoken.

This is the prize: not a fortress, but a country that fails well — one in which an attacker pays more than the result is worth, and a citizen, watching the system bend, still trusts that it will hold.

Choose what must not fail. Decide where trust comes from.

Raise the floor. Defend at machine speed. Mobilise the country.

Tell the truth first.



Bibliography

Threat intelligence and AI-enabled cyber misuse

Anthropic. (2025, November 13). Disrupting the first reported AI-orchestrated cyber-espionage campaign. <https://www.anthropic.com/news/disrupting-AI-espionage>

Google Cloud / Mandiant. (2026). M-Trends 2026: Data, insights, and strategies from the frontlines. Google Cloud Blog. <https://cloud.google.com/security/resources/m-trends>

Google Cloud / Mandiant. (2024, October). How low can you go? An analysis of 2023 time-to-exploit trends. Google Cloud Blog. <https://cloud.google.com/blog/topics/threat-intelligence/time-to-exploit-trends-2023>

Google Threat Intelligence Group. (2026, March 5). Look what you made us patch: 2025 zero-days in review. Google Cloud Blog. <https://cloud.google.com/blog/topics/threat-intelligence/2025-zero-day-review>

Zero Day Clock. (n.d.). Time-to-exploit visualisation, 2018–2026. <https://www.zerodayclock.com/>

AI risk and frontier capability

Anthropic. (2025). Responsible Scaling Policy (current version). <https://www.anthropic.com/responsible-scaling-policy>

Google DeepMind. (2025). Frontier Safety Framework. <https://deepmind.google/discover/blog/introducing-the-frontier-safety-framework/>

National Institute of Standards and Technology. (2025). Adversarial machine learning: A taxonomy and terminology of attacks and mitigations. NIST AI 100-2 E2025. <https://doi.org/10.6028/NIST.AI.100-2e2025>

UK AI Security Institute. (n.d.). Frontier AI evaluations and cyber-risk research. Government of the United Kingdom. <https://www.aisi.gov.uk/>

UK National Cyber Security Centre. (2026). Why cyber defenders need to be ready for frontier AI; Preparing for a vulnerability patch wave; 10

questions to ask when using AI models to find vulnerabilities. <https://www.ncsc.gov.uk/>

International Monetary Fund. (2026). Financial stability risks mount as artificial intelligence fuels cyberattacks. <https://www.imf.org/>

Elderson, F. (2026). Interview on AI and operational resilience in banking. European Central Bank. <https://www.ecb.europa.eu/>

Digital identity, cryptography and trust services

European Parliament and Council. (2024). Regulation (EU) 2024/1183 amending Regulation (EU) No 910/2014 (eIDAS) — European Digital Identity Framework. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1183/oj/eng>

Information System Authority (RIA). (n.d.). Electronic identification and trust services. Republic of Estonia. <https://www.ria.ee/en/state-information-system/eid>

National Cyber Security Centre. (2025, March 20). Timelines for migration to post-quantum cryptography. Government of the United Kingdom. <https://www.ncsc.gov.uk/guidance/pqc-migration-timelines>

National Institute of Standards and Technology. (2024). Post-quantum cryptography standards: FIPS 203, FIPS 204, FIPS 205. U.S. Department of Commerce. <https://csrc.nist.gov/projects/post-quantum-cryptography>

X-Road. (n.d.). X-Road: Data exchange layer for digital societies. Nordic Institute for Interoperability Solutions. <https://x-road.global/>

Regulation and operational resilience

European Parliament and Council. EU cyber and AI regulations: NIS2 — Directive (EU) 2022/2555; DORA — Regulation (EU) 2022/2554; Cyber Resilience Act — Regulation (EU) 2024/2847; AI Act — Regulation (EU) 2024/1689; Cyber Solidarity Act — Regulation (EU) 2025/38. Official Journal of the European Union. <https://eur-lex.europa.eu>

European Union Agency for Cybersecurity (ENISA). (n.d.). EU-CyCLONE — European cyber crisis liaison organisation network. <https://www.enisa.europa.eu/topics/incident-response/cyclone>

Schmitt, M. N. (Ed.) (2017). Tallinn Manual 2.0 on the international law applicable to cyber operations. Cambridge University Press.

North Atlantic Treaty Organization. (2021, June 14). Brussels Summit Communiqué, paragraph 32. https://www.nato.int/cps/en/natohq/news_185000.htm

Digital state resilience, historical comparators and crisis response

Raieste, A., Rebane, A., Tapupere, M., & McBride, K. (2024). Government Resilience in the Digital Age. Oxford Internet Institute, Estonian IT Centre and Nortal. Tallinn.

National Audit Office. (2017, October 27). Investigation: WannaCry cyber attack and the NHS. UK National Audit Office. <https://www.nao.org.uk/reports/investigation-wannacry-cyber-attack-and-the-nhs/>

Ottis, R. (2008). Analysis of the 2007 cyber attacks against Estonia from the information warfare perspective. NATO Cooperative Cyber Defence Centre of Excellence. <https://ccdcoe.org/>

