SOVEREIGN DIGITAL RESILIENCE · PAPER 3 OF 3

Exit-Capable by Design

Vendor Concentration Is a Sovereignty Risk

SANJEET KUMAR

Enterprise Architecture &
Technology Governance

Critical Infrastructure
Sovereign Digital Resilience
AI Governance

sanjeetkumar.com

Exit-Capable by Design: Vendor Concentration Is a Sovereignty Risk

Sanjeet Kumar — Enterprise Architecture & Technology Governance Leader *Critical Infrastructure · Sovereign Digital Resilience · AI Governance — July 2026. Views are the author's own.*

ORCID 0009-0003-1677-684X · DOI: 10.5281/zenodo.21970869

All deposits and versions: sanjeetkumar.com/papers/exit-capable-by-design

Two decades ago The Open Group's Jericho Forum argued for "de-perimeterisation" — the then-heretical idea that the network boundary was dissolving and security had to move to the data and the transaction. The industry eventually renamed that idea Zero Trust. I've carried one conviction from those arguments ever since: **the perimeter you should worry about is never only the network one.**

Today the perimeter quietly dissolving is *commercial*. Critical systems — registries, payment rails, government platforms — increasingly depend on a small number of hyperscale and specialist vendors, and that concentration is a sovereignty problem wearing a procurement badge. A jurisdiction whose system of record can only run, or only be understood, inside one supplier's ecosystem has ceded a measure of control no data-residency clause can recover. Foreign-access laws, licensing changes, product sunsets, acquisitions, sanctions regimes: any of these can reach your crown jewels through the vendor relationship, entirely legally, without a single packet crossing your firewall.

The regulators arrived, and they disagree about the remedy

When I first set this argument down, exit-capability was an architectural preference. It is now, across four jurisdictions, a supervisory expectation — and the interesting part is not that the regulators agree with the diagnosis. It is that they have reached for three different remedies.

Oblige the institution. Canada's OSFI, in its revised Guideline B-10, splits concentration risk in two: institution-specific risk, where one institution over-relies on a single third party, subcontractor or geography across multiple activities; and systemic risk, where one provider or one geography serves many institutions at once. The distinction is worth borrowing whatever sector you are in, because only the first is inside your control and only the second explains why your own diversification may not help you. B-10 requires exit and transition-out planning across the third-party lifecycle, with exit strategies expected for cloud arrangements assessed as high-risk or critical. The EU's DORA takes the same route in Article 28: a register of ICT arrangements, pre-contractual due diligence, concentration assessment, and maintained exit strategies for every arrangement supporting a critical or important function. The Central Bank of the UAE arrived earlier and from a different direction, through its Outsourcing Regulation and Standards for Banks, which hold that the bank retains ownership of its data, that the regulator may access it on request, and that offshore arrangements carry additional restrictions.

Supervise the provider. The second remedy concedes that firm-level diligence cannot reach a supplier serving the whole market. DORA's Articles 31 to 44 place designated Critical ICT Third-Party Providers under direct European oversight, with substitutability among the designation criteria. The United Kingdom built the same instrument under the Financial Services and Markets Act: the critical third parties rules in PS16/24 took effect on 1 January 2025 and then sat dormant, applying to nobody, until HM Treasury made the first designations. Two features of that regime deserve attention from anyone drafting a board paper. Designation is not authorisation — an overseen provider has not been vouched for. And designation does not reduce the obligations of the firms relying on it: the register, the concentration assessment and the exit plan remain yours.

Make the market switchable. The third remedy is the only one that changes the economics rather than the paperwork. Chapter VI of the EU Data Act, applicable since 12 September 2025, requires providers of data processing services to remove the contractual, technical and commercial obstacles to switching: a notice period capped at two months, a transitional period of thirty days extendable where migration is technically unfeasible, an exhaustive specification of what can be ported, and — from 12 January 2027 — the elimination of switching and egress charges. It applies to providers serving EU customers wherever they are established.

For anyone running a critical system the practical consequence is the same under all three: you must be able to leave, and you must be able to show it. What has changed is the standard of proof. An untested exit plan used to be a governance weakness that a determined architect could raise and lose. It is now a finding.

The architectural response I practice and advocate is **exit-capability as a design requirement** — resilience expressed as the permanent, tested ability to leave:

Decouple the record of authority from the platform of convenience. Authoritative data lives in open, documented formats; the vendor's value-add operates on top of it, never as the sole means of interpreting it. If your system of record is only legible through the vendor's software, the vendor owns your record in every way that matters.

Make the exit path a maintained artefact. An exit plan written at contract signing and never touched is theatre. Exit-capability means the migration path is documented, cost-estimated, and periodically exercised — treated like disaster recovery, because that's what it is: recovery from a *commercial* disaster.

Measure concentration deliberately. A technology heatmap that maps capabilities to platforms makes concentration visible and disciplinable: which critical functions share one vendor, one region, one licensing regime? Multi-year roadmaps and disciplined decommissioning then become the instruments that *reduce* the number — modular design, one retirement at a time. Simplification is not cost-cutting; it is de-risking with a budget line.

Extend the lens to AI immediately. Model dependencies are the fastest-forming concentration risk in the enterprise today. An organization whose workflows assume one proprietary model family — its behaviours, context formats, and pricing — is rebuilding the vendor lock-in of the 1990s at the cognition layer. The same disciplines apply: abstraction at the

interface, portable context, and a tested answer to "what if we had to switch models in a quarter?"

The model layer is not covered

The switching regimes above govern data processing services. They do not govern model behaviour, and the distinction is about to matter a great deal. You can port every byte you hold and still be unable to reproduce your system, because what you depended on was not the data but how a particular model family responded to it — its refusals, its formats, its tolerance for a prompt written three years ago by someone who has left. Portability at the cognition layer means holding the artefacts that let you re-qualify a replacement: your evaluation sets, your acceptance thresholds, your context in a form that is not one vendor's schema. Treat the model as a replaceable component and the evaluations as the thing you own. Anything else rebuilds the lock-in of the 1990s one layer up, and this time the dependency is on behaviour nobody can specify.

What an exit plan must contain to be more than theatre

Most exit plans I have reviewed fail on the same five points. They are not difficult; they are simply unglamorous, and nobody is promoted for finishing them.

A named alternative. "We would go to market" is not a plan; it is a description of not having one. Name the destination — a specific competing platform, an open-source stack, or the decision to bring the function in-house — and accept that naming it will be uncomfortable, because it forces you to admit what the alternative would cost you in capability.

A tested extraction path. Not a contractual right to your data: a rehearsed export, with the format recorded, the volume recorded, the elapsed time recorded, and the name of whoever last ran it. Contractual portability that has never been exercised is an assertion, and assertions are what this whole argument is about.

The interpretation layer. This is the one almost everybody omits. Your data leaves in an open format and your business rules stay behind — in workflow configuration, in report definitions, in field-level validation, in the twenty-year accretion of settings nobody documented because the platform held them. Data without its interpretation is an archive, not a system. Ask what fraction of your operating logic lives in configuration rather than in code you hold, and treat the answer as a lock-in metric.

A refreshed cost and time estimate. Estimates age faster than architectures. A migration costed at contract signature and never revisited tells you what a different system would have cost a different organisation. Refresh annually, alongside the disaster-recovery test, because that is what this is: recovery from a commercial disaster.

Trigger conditions and a named decision-maker. Write down what would make you go — an acquisition, a licensing change, a sustained breach of service levels, a jurisdictional development — and who decides. Exit decisions taken for the first time during a crisis are taken badly, and usually too late to matter.

Concentration hides one layer down

A heatmap that counts vendors will understate your exposure, because concentration accumulates along dimensions that procurement does not track. Count instead by region, by licensing regime, by identity plane — the single sign-on estate is the most under-examined concentration in most enterprises — and by model family. Then count your fourth parties: two suppliers you chose deliberately for diversity may share one subcontractor, one region, or one upstream dependency, and you will not discover it from your own contracts. This is precisely why both OSFI and DORA push institutions toward a register rather than a policy. The inventory is the instrument; the policy is a preface to it.

None of this is anti-vendor. Great suppliers are leverage; the objective is a relationship where *both* parties know you could leave — which, not incidentally, is also when pricing and service are at their best. Sovereignty was never about isolation. It is about retaining the option to choose, at every layer: network, data, jurisdiction — and vendor.

De-perimeterisation taught us to stop trusting the network boundary. Its modern corollary: stop trusting the commercial one.

Five questions for a board

1. For each critical system, who is the named alternative provider, and when was the extraction path last tested end to end?

2. What proportion of our operating logic lives in vendor configuration rather than in artefacts we hold and can read without the vendor?

3. Which of our critical functions share a single vendor, region, licensing regime or identity plane — and do any of our diversified suppliers share a fourth party?

4. What would a forced migration of our most concentrated dependency cost, in money and elapsed time, on an estimate refreshed within the last twelve months?

5. If our principal AI provider changed its pricing, its terms or its model behaviour next quarter, what would we do, and who is authorised to decide it?

References

1. Office of the Superintendent of Financial Institutions (Canada), Guideline B-10: Third-Party Risk Management, 2023.

2. Regulation (EU) 2022/2554, Digital Operational Resilience Act, Chapter V, Articles 28–44; applicable from 17 January 2025.

3. Bank of England, Prudential Regulation Authority and Financial Conduct Authority, PS16/24 and PS24/16, Operational resilience: Critical third parties to the UK financial sector, 12 November 2024; rules in force 1 January 2025.

4. Bank of England, SS6/24, Critical third parties to the UK financial sector.

5. Regulation (EU) 2023/2854, Data Act, Chapter VI, Articles 23–31; applicable from 12 September 2025, with switching and egress charges eliminated from 12 January 2027.

6. Central Bank of the United Arab Emirates, Outsourcing Regulation for Banks and Outsourcing Standards for Banks, Circular No. 14/2021.

7. The Open Group, Jericho Forum Commandments, 2007.

8. National Institute of Standards and Technology, SP 800-207, Zero Trust Architecture, 2020.

9. Kumar, S., Century-Scale Systems: What Land Registries Teach Us About Data Sovereignty, Zenodo, 2026. doi:10.5281/zenodo.21970840

10. Kumar, S., Governing Agentic AI: In Systems of Public Record — Governance Is the Accelerant, Zenodo, 2026. doi:10.5281/zenodo.21971009

11. Kumar, S., Decision Custody: Who Answers for a Decision After the Machine That Made It Is Gone, Zenodo, 2026. doi:10.5281/zenodo.21959271

12. Kumar, S., The Cold Standard — A Fixed Reference for a Record That Won't Stop Moving, Zenodo, 2026. doi:10.5281/zenodo.22073326

AI-assistance disclosure: This paper was drafted and refined with AI assistance. The arguments, positions and judgements are the author's own, and the regulatory sources cited above were verified against their primary records before publication.