Centre for **Risk Studies**

UNIVERSITY OF CAMBRIDGE Judge Business School

McKinsey & Company

# Cambridge - McKinsey Risk Prize Bio-sketch and Photo Page

Photo

**Student name:** Meheer Commuri

**Email contact:** mlc202@cam.ac.uk

**Date of submission:** March 31, 2026

**Title of submission:** The State as Tenant: Cloud Concentration and the Sovereign Risk No One Is Pricing

**I am a candidate for the degree:** Master of Philosophy

**Department:** Politics and International Studies

**LinkedIn profile link:** https://www.linkedin.com/in/meheer-commuri/

**Biographical Sketch:** Meheer Commuri is an M.Phil student in Politics and International Studies at the University of Cambridge. His research investigates how states build, sustain, and lose competitive advantage in an era when the determinants of power extend well beyond military capacity into infrastructure, technology, and institutional design. His current work examines how regional powers arbitrage among competing partnerships offered by major powers, and his previous research explored how authoritarian regimes rehabilitate their standing in the international community. He brings formal models and quantitative methods to these questions. Before Cambridge, Meheer completed his BA with Distinction, *summa cum laude*, in International Relations at the University of Pennsylvania, where he won the best thesis award in International Security. He holds dual minors in Political Science and Survey Research and Data Analytics.

Centre for
**Risk Studies**

UNIVERSITY OF CAMBRIDGE
Judge Business School

McKinsey & Company

# Cambridge - McKinsey Risk Prize

# Declaration Form

**Student name:** Meheer Commuri

**Email contact:** mlc202@cam.ac.uk

**Date of submission:** March 31, 2025

**Title of submission:** The State as Tenant: Cloud Concentration and the Sovereign Risk No One Is Pricing

**Number of words of submission:** 4,100 total (3,770 essay, 330 references).

**I am a candidate for the degree:** Master of Philosophy

**Academic Institution/Department:** Politics and International Studies

**Declaration**

I confirm that this piece of work is my own and does not violate the University of Cambridge Judge Business School's guidelines on Plagiarism.

I agree that my submission will be available as an internal document for members of both Cambridge Judge Business School and McKinsey & Co's Global Risk Practice.

If my submission either wins or receives an honourable mention for the Risk Prize, then I agree that (a) I will provide a recorded 2 minute overview of my paper, (b) my submission can be made public on a Cambridge Judge Business School and/or McKinsey & Co websites.

This submission on risk management does not exceed 10 pages.

Signed (Electronic Signature)

*Meheer Commuri*

Please include this declaration form after the cover page of your paper submission.

# The State as Tenant: Cloud Concentration and the Sovereign Risk No One Is Pricing

Meheer Commuri

March 2026

## 1. Introduction

On 25 October 2025, a routine fault in Amazon Web Services' US-East-1 region simultaneously disabled Lloyds Banking Group's payment processing, delayed London Stock Exchange settlement, and rendered HMRC's digital tax portal inaccessible for fifteen hours. No adversary was involved, and no single organisation had been negligent; a technical malfunction in one company's Virginia data centre had temporarily suspended core functions of a sovereign state. The incident was not unusual, as it was the seventh major cloud outage in five years to disrupt critical public infrastructure, and none of the prior events had prompted structural reform. However, what the pattern revealed was a category of risk that existing frameworks do not recognise: the exposure when sovereign capacity becomes contingent on privately controlled, highly concentrated infrastructure.

This essay argues that cloud infrastructure concentration has produced a new form of sovereign-systemic risk. As governments have migrated critical functions onto a cloud market dominated by three providers controlling 63 per cent of global capacity, they have, in a structural sense, become tenants of the infrastructure through which they govern. States are not buying computing; they are leasing state capacity from private firms. The resulting exposure operates on three levels. At the first, cloud concentration is an economic risk to firms and GDP. At the second, it is a risk to the delivery of public functions on which citizens depend. At the third, it is a risk to the legitimacy and operational sovereignty of the state itself, because it produces a condition in which governments retain formal authority but lack the infrastructural capacity to exercise or protect that authority when disruption occurs, and sovereignty itself risks becoming vestigial. The exposure at every level is substantial, it is compounding, and it is almost entirely uninsured.

Drawing on outage-by-outage loss estimates from the Parametrix Cloud Monitoring System, market concentration data from Synergy Research Group, and insurance gap analysis from the Geneva Association and Guy Carpenter, we first quantify this exposure. We then identify why three independent failure modes in risk modelling, insurance markets, and government oversight allow this exposure to persist unmanaged. Finally we propose structural alternatives including a diagnostic index that any state or advisory body can apply to calibrate the potency of this risk.

## 2. Quantifying the Exposure

Recent cloud outages have revealed that the relationship between cloud market concentration and outage severity is more structural than coincidental. According to Synergy Research Group, the combined market share of AWS, Microsoft Azure, and Google Cloud Platform rose from approximately 58 per cent in 2020 to 63 per cent by 2025. Over the same period, annual significant outage hours tracked by the Parametrix Cloud Monitoring System increased from approximately 80 to over 150, with the economic impact per event rising disproportionately as more critical services concentrated on fewer platforms (Parametrix, 2024a; 2025a; 2025b). Exhibit 1 traces both trends together, and it is evident that the upward slopes are mutually reinforcing, where higher concentration enlarges the blast radius of each incident, which in turn amplifies the economic

severity of each additional outage hour.

**Exhibit 1. Rising Cloud Outage Severity Amid Persistent Market Concentration**

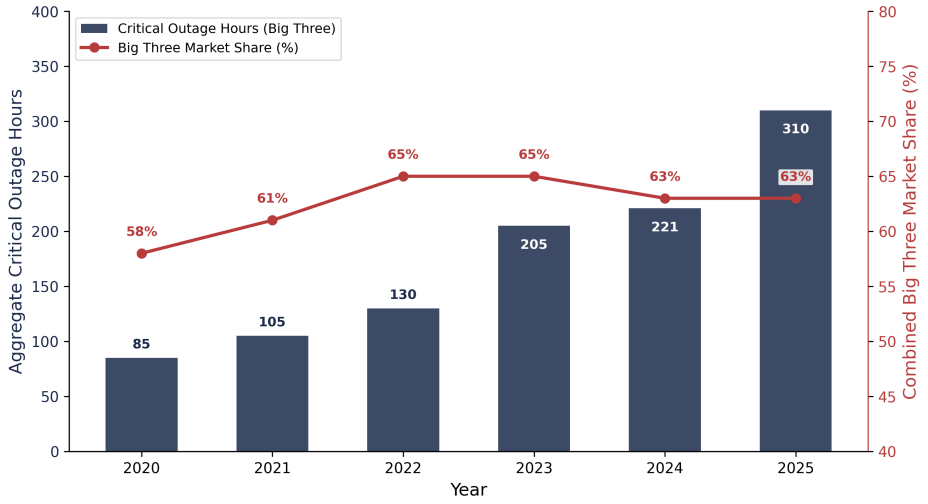


*Sources: Parametrix Cloud Outage Risk Reports (2022–2024); Synergy Research Group (2020–2025 Q3). 2020–2021 outage hours estimated from Uptime Institute benchmarks; 2025 includes Oct AWS & Azure events.*

But beneath this trend lies a reality whose implications are more urgent than we realise. It reveals that a system can be profitable, innovative, and still socially underinsured with crippling consequences. Exhibit 2 presents seven major cloud outage events between 2020 and 2025, comparing estimated total economic losses against insured losses. Across every event, the uninsured proportion ranges from 85 to 95 per cent, with the weighted average protection gap sitting at approximately 90 per cent.

**Exhibit 2. The Cyber Protection Gap: Economic Losses vs. Insured Coverage for Major Cloud Outage Events (2020–2025)**

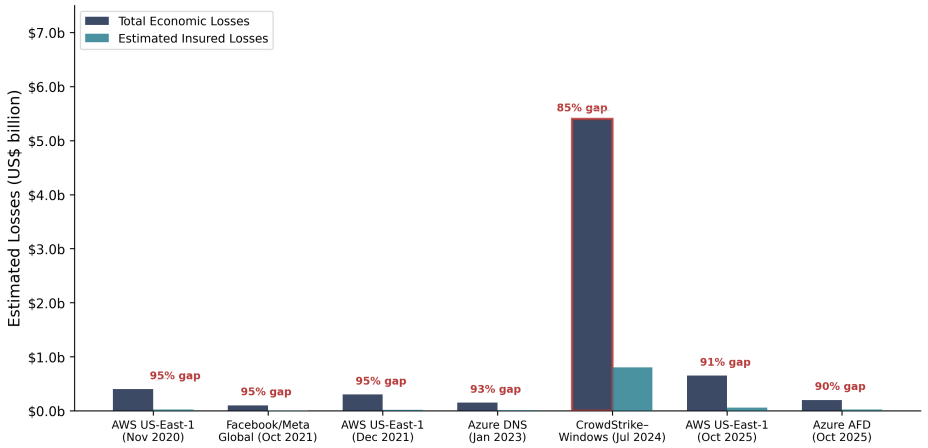


*Sources: Parametrix event analyses (2024b, 2025b, 2025c); Fortune/Axios (CrowdStrike); Geneva Association (2024). Azure Oct 2025 estimated at one-third of AWS Oct 2025 per Parametrix comparative footprint analysis. Insured losses estimated at 5–20% of economic losses. Geneva Association estimates aggregate cyber protection gap at ~90%.*

The CrowdStrike-Windows incident of July 2024 is perhaps one of the most concrete and recent demonstrations of this structural underinsurance against a significant risk. A faulty software update distributed through CrowdStrike's Falcon platform disabled an estimated 8.5 million Windows devices globally, grounding flights, freezing hospital systems, and halting financial settlement (Parametrix, 2024b). Total economic losses reached approximately $5.4 billion, whereas insured claims totalled roughly $800 million (Guy Carpenter, 2025). The remaining $4.6 billion was absorbed by affected organisations as uncompensated loss. In other words, infrastructure risk sits lowest on the balance sheet of the provider, and highest on the balance sheet of the society that depends on it.

This protection gap is not merely a reflection of market immaturity, but it is more of a fundamental

incompatibility between the characteristics of cloud concentration risk and the requirements of traditional insurance. Diversification, the actuarial foundation of insurability, assumes that insured risks are largely independent of one another. But when three providers underpin the majority of global cloud workloads, a single failure can simultaneously affect millions of policyholders. The correlated nature of the exposure places it closer to catastrophe risk than to conventional operational loss, yet no dedicated catastrophe-style risk transfer mechanism currently exists for cloud concentration risk (Geneva Association, 2024).

A puzzling pricing paradox sharpens this concern because, despite rising frequency and severity of cloud-related incidents, US cyber insurance premiums declined 4 per cent in Q1 2025 (Marsh McLennan, 2025). Competitive pressure among insurers, combined with favourable loss ratios in recent quarters, has driven rates downward at precisely the moment when systemic exposure has been growing. This divergence between pricing trajectory and underlying risk trajectory is characteristic of a mispriced systemic exposure, and we argue in this essay (Section 5) that it has identifiable structural causes.

## 3. From Operational Risk to Sovereign Risk

Standard risk taxonomies classify cloud outages as operational disruptions affecting individual firms or, at most, as supply-chain risk. But this classification misidentifies the nature of the exposure because the portfolio of tenants for cloud computing is distinct and unprecedented. When government functions depend on infrastructure controlled by a small number of private providers, an outage does not merely interrupt commercial operations; it suspends the capacity of the state to discharge its obligations to citizens. This is a distinct form of risk, and it is without precedent in the present digital form.

Research on how infrastructure shapes institutional power (Star and Ruhleder, 1996; Plantin et al., 2018; Mann, 1984) provides a useful framework here. The state is not a customer of cloud services in the ordinary commercial sense but a tenant whose operational capacity depends on infrastructure it does not own, cannot inspect, and cannot independently restore, and tenancy carries a risk that procurement does not, because a tenant's access is always conditional on the provider's continued willingness and capacity to serve.

We map in Exhibit 3 government functions along two dimensions: the degree of cloud dependency and the level of government resilience and redundancy for that function. Functions in the lower-right quadrant, which include tax collection and revenue, benefits and welfare administration, health records, financial regulation, and emergency services, combine high dependency on cloud infrastructure with low government resilience, placing them in the zone of critical sovereign exposure. When these functions fail, the consequences extend well beyond economic loss, because citizens cannot access entitlements to which they hold legal right, patients cannot retrieve medical records essential to continuity of care, and tax authorities cannot process the obligations that fund the apparatus of the state. For the duration of such an outage, the state's capacity to fulfil its most basic compact with citizens is effectively suspended.

The mapping in Exhibit 3 reveals a condition in which the formal apparatus of state authority persists, but the infrastructure required to exercise that authority has migrated beyond sovereign control, and sovereignty becomes vestigial, retained in name but hollowed of the capacity that once gave it meaning. This condition differs from conventional outsourcing in a critical respect, because in standard outsourcing arrangements, governments contract for specific services while retaining the practical option to bring functions in-house or switch providers. Infrastructural dependency

operates differently, more insidiously, and carries a different order of risk. Technical expertise, data architectures, and operational workflows become so deeply integrated with specific platforms that switching costs grow prohibitive over time, path dependencies accumulate, and what looks like outsourcing quietly becomes constitutional drift. Governments discover not that they have contracted out a service, but that they have transferred a capacity.

**Exhibit 3. Sovereign Function Vulnerability Map: Cloud Dependency vs. Government Resilience**

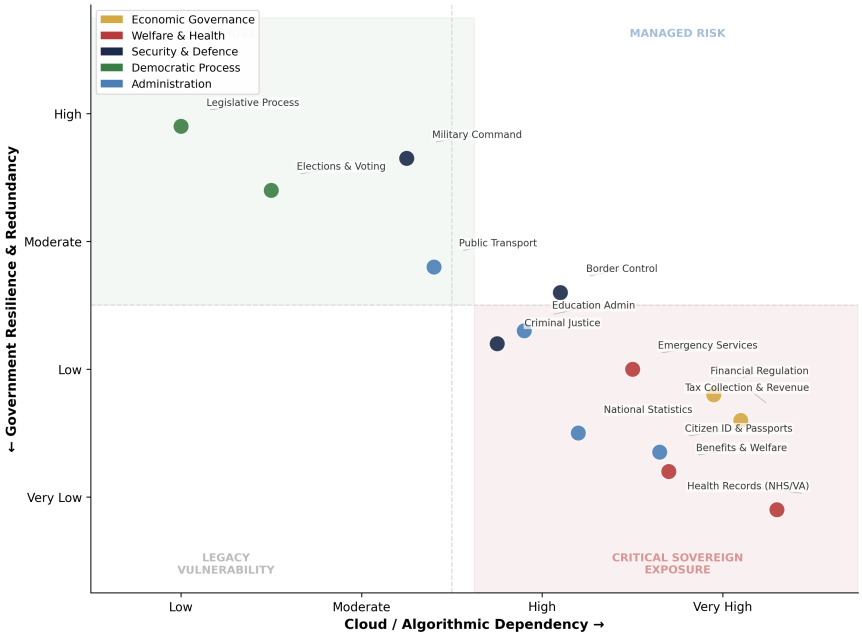


*Sources: Author analysis drawing on Table 4 of Commuri (2026); OECD (2024); UCL (2025); National Security Commission on AI (2024). Positions are illustrative estimates based on publicly available assessments of government cloud adoption and contingency planning.*

When a cloud outage suspends government functions, the affected population experiences a withdrawal of state capacity that no elected official authorised and no democratic process can remedy in real time. The more invisible the infrastructure becomes, the more consequential its governance becomes, because its defaults, its points of failure, and its terms of service effectively set the operational boundaries of public authority.

## 4. Quantifying the Sovereign Risk

To project the trajectory of this exposure, we construct a GDP-at-risk estimate anchored on Parametrix's empirical data (Exhibit 4). The projections that follow are structured diagnostics rather than point forecasts; their value lies in establishing orders of magnitude and directional trends rather than precise annual figures. The October 2025 AWS outage, which lasted approximately 15 hours and affected an estimated 30 per cent of monitored cloud providers' downstream services, generated losses in the range of $500–650 million (Parametrix, 2025b), implying a direct loss rate of approximately $37–43 million per hour. When we include indirect effects such as supply-chain delays, lost productivity, and reputational harm, the effective rate could rise to approximately $55 million per hour.

This rate reflects 2025 conditions, in which approximately 20 per cent of US GDP depended on cloud infrastructure. That economic exposure is relevant here not because sovereign risk is reducible to commercial loss, but because it indicates the scale of societal and administrative activity now contingent on digital infrastructure concentrated in a small number of private providers. As cloud adoption accelerates, so does the share of economic output and public-service delivery exposed to concentration risk. We therefore project forward to 2030 using three inputs: US GDP growth from

the IMF World Economic Outlook (2025), cloud adoption share from Gartner (2024) and Synergy Research Group, and a vulnerability growth factor reflecting deepening integration and diminishing substitutability as path dependencies compound.

**Exhibit 4. The Sovereign Risk Multiplier:**
**Estimated US GDP-at-Risk per Hour of Hyperscaler Downtime (2020–2025)**

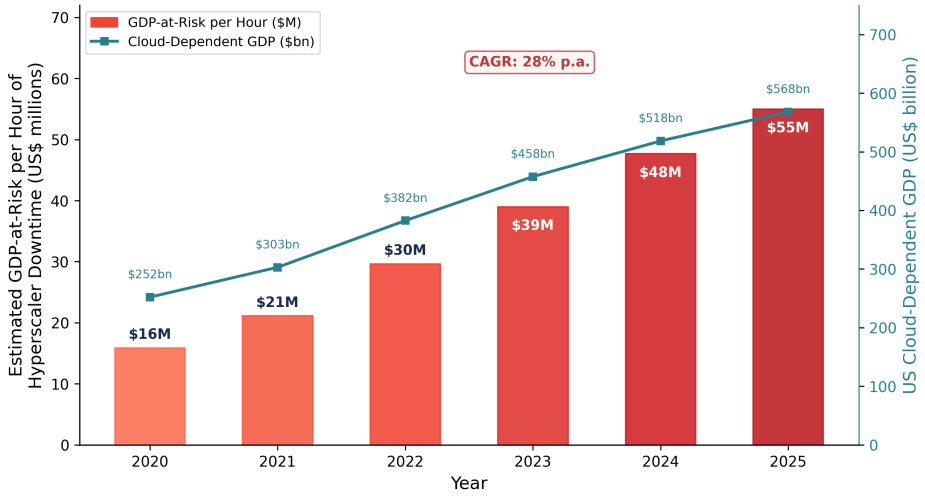


*Sources: IMF WEO (US GDP); Katz/TAS (2024) cloud GDP contribution study; Parametrix (Oct 2025 AWS outage loss estimates: $500–650M over 15 hours). Methodology: Empirically anchored on Parametrix 2025 estimate, scaled by cloud-dependent GDP and integration depth. Excludes cascade effects, reputational losses, and sovereign capacity costs.*

On this basis, annual GDP-at-risk from major cloud outages rises from approximately $20 billion in 2020 to $55 billion in 2025 and is projected to reach $115–140 billion by 2030 under current concentration trends. These estimates are not forecasts of realised annual losses but a proxy for the scale of activity whose continuity depends on the uninterrupted functioning of concentrated cloud infrastructure beyond direct sovereign control. For the purposes of this essay, that matters because the same dependence that places output at risk also places state capacity at risk: the ability to make payments, administer benefits, collect revenue, coordinate emergency response, and maintain continuity in core public functions. The sovereign risk multiplier, which we define as the ratio of GDP-at-risk in year $t$ to the 2020 baseline, therefore measures not merely growing commercial exposure but the expanding reach of cloud concentration for the state. The compounding trajectory reflects not only increased cloud adoption but deepening integration, in which proliferating interdependencies cause the blast radius of any single outage to scale non-linearly with the density of the system it disrupts.

### 4.1 The Sovereign Infrastructure Exposure Index

To make this analysis operational, we propose the Sovereign Infrastructure Exposure Index (SIE), a diagnostic that any state or advisory body can apply. The index combines three risk inputs, each scaled from 0 to 1: **D** (sovereignty exposure), the proportion of legitimacy-critical government functions running on commercial cloud infrastructure; **C** (concentration risk), the market share of the single largest provider in the government's cloud estate; and **L** (lock-in risk, defined as $1 - S$, where $S$ represents demonstrated substitutability), measuring how much option value the state has surrendered. The index is multiplicative:

$$\text{SIE} = D \times C \times L \tag{1}$$

This structure encodes a specific analytical claim, which is that sovereign infrastructure risk arises from the interaction of dependency, concentration, and lock-in, and the presence of any one factor alone does not generate dangerous exposure. A state that hosts most functions on a single provider

but maintains genuine exit capacity (low L) faces a categorically different risk profile from one that appears diversified on paper but could not actually migrate under duress (high L). What the index penalises is their combination, which is precisely the structural condition that converts a routine outage into a sovereign event.

We propose four categories of exposure: Category 1 (SIE below 0.10), resilient, where sovereign functions are diversified or self-hosted; Category 2 (SIE 0.10–0.25), exposed but manageable through targeted multi-cloud mandates or portability investments; Category 3 (SIE 0.25–0.50), structurally vulnerable, where path dependencies are already constraining available options; and Category 4 (SIE above 0.50), critical dependency, where sovereign capacity is contingent on a single provider's stability and remediation is both urgent and costly.

Table 1 maps six countries against the index, and the results illustrate a counterintuitive pattern. Conventional rankings of digital government maturity tend to focus on deployment while underplaying the risk, and would place Denmark well above countries such as India. Saudi Arabia would scarcely appear in most discussions of digital governance leadership. Yet the order inverts on sovereign infrastructure exposure. India's MeghRaj Government Cloud and the India Stack architecture provide genuine fallback capacity that the country's level of digitisation does not strictly demand, but that its sovereignty strategy deliberately built. Saudi Arabia, through its NEOM-linked cloud investments and mandatory data localisation under the National Data Management Office, has constructed substitutability that wealthier per-capita states in Europe have not. France benefits from OVHcloud and its *cloud de confiance* policy, which created domestic alternatives that a comparably sized economy like the United Kingdom has not developed. Denmark, celebrated as one of Europe's most advanced digital governments, scores near the top of Category 3 precisely because its depth of digitisation through US cloud providers produced the highest dependency and the lowest substitutability in the sample, while Thailand, despite a less mature digital government overall, scores similarly because its rapid cloud adoption concentrated on a single provider with negligible sovereign fallback. The comparison across Table 1 reveals that digital advancement and digital resilience are measuring fundamentally different things, and what conflates them is itself a source of risk.

**Table 1:** Illustrative SIE Scores for Selected Jurisdictions

| **Jurisdiction** | **D** | **C** | **L** | **SIE** | **Category** |
|---|---|---|---|---|---|
| India | 0.35 | 0.35 | 0.45 | 0.06 | 1 (Resilient) |
| Saudi Arabia | 0.40 | 0.45 | 0.50 | 0.09 | 1 (Resilient) |
| France | 0.55 | 0.40 | 0.55 | 0.12 | 2 (Manageable) |
| United Kingdom | 0.65 | 0.55 | 0.75 | 0.27 | 3 (Vulnerable) |
| Thailand | 0.55 | 0.65 | 0.90 | 0.32 | 3 (Vulnerable) |
| Denmark | 0.75 | 0.60 | 0.85 | 0.38 | 3 (Vulnerable) |

*Note: Scores are indicative estimates derived from published procurement data, digital government surveys (OECD, 2024), and publicly reported cloud strategies. Precise calibration would require jurisdiction-specific audit data.*

The SIE index also makes legible the effect of specific interventions. For example, mandating multi-cloud deployment for critical functions directly reduces C; investing in sovereign capacity for legitimacy-sensitive services reduces D; and maintaining portability standards together with regularly tested exit procedures lowers L. A state in Category 3 that meaningfully improves any one input can shift toward Category 2, but a state in Category 4, where all three risk factors are elevated, must act on at least two fronts simultaneously, because the multiplicative structure implies that a single intervention cannot compensate when dependency, concentration, and lock-in

are all high.

### 4.2 The Cost of Inaction

The static scores in Table 1 describe the present, yet risk is dynamic. Exhibit 5 projects SIE trajectories for each jurisdiction through 2030, assuming a continuation of current policy settings and adoption trends. What emerges from the projection is a set of paths that diverge far more sharply than the current scores alone would suggest. India's score declines slightly as sovereign infrastructure investments mature and substitutability improves. Saudi Arabia and France remain broadly stable, because their earlier policy choices established structural floors beneath their exposure. The United Kingdom rises steadily as digitisation deepens without compensating investment in sovereign capacity.

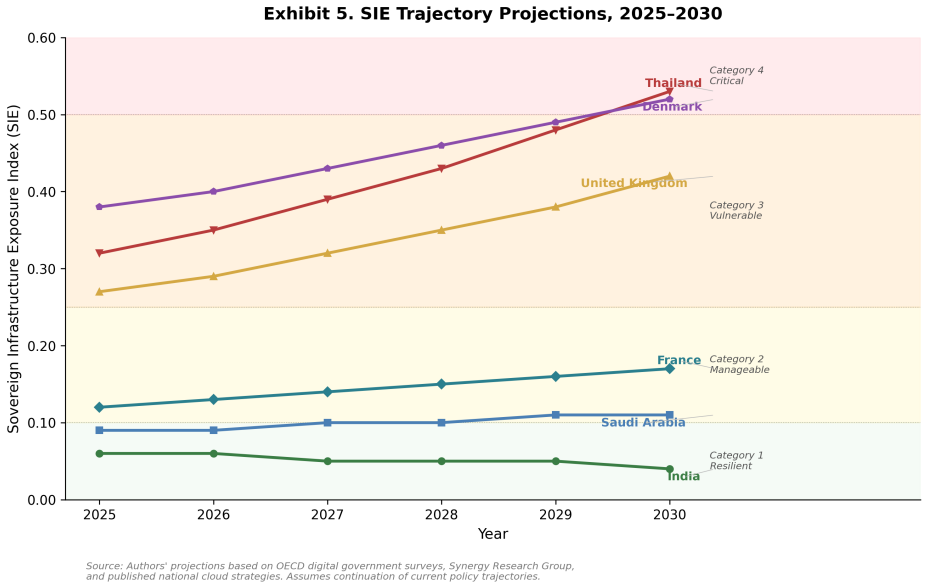


The most consequential trajectories, however, belong to Denmark and Thailand, both of which cross the Category 4 threshold by 2030 on current trends. Denmark's trajectory is particularly instructive because it illustrates how a state can move from celebrated digital leader to critical sovereign dependency without any single decision appearing reckless. Each incremental step toward deeper cloud integration was individually rational; their cumulative effect is a structural exposure that will be difficult and expensive to reverse. Thailand's trajectory follows a steeper path to the same destination, driven by rapid adoption on a single provider with no parallel investment in exit capacity.

Exhibit 5 also exposes the asymmetry of action and inaction, and the corresponding hysteresis. The states that invested early in substitutability (India, Saudi Arabia, France) face a flat or declining risk trajectory at modest ongoing cost. The states that did not are now on compounding curves where the cost of remediation rises with each year of delay. This asymmetry is the strongest argument for early intervention, because the same multi-cloud mandate that would cost a Category 2 state relatively little to implement becomes a disruptive and expensive restructuring exercise for a state that has allowed itself to drift into Category 4.

## 5. How We Got Here, and What Could Work

Three interlocking failures explain why this risk persists without systematic management, and we review each in turn before proposing structural remedies.

**Risk models miss the correlation.** Standard enterprise risk management treats cloud dependency

as a supplier risk for individual organisations, an approach that is structurally incapable of capturing the systemic dimension of the exposure. Recent work on critical infrastructure resilience has recognised that single-vendor failure can constitute a structural vulnerability, and that cascading effects through fourth- and fifth-party dependencies amplify disruption further than firm-level assessments anticipate (McKinsey & Company, 2024). This observation is correct as far as it goes, but it remains framed at the level of the individual organisation. When millions of organisations depend on the same three providers, what presents as idiosyncratic risk at the firm level aggregates into heavily correlated exposure at the system level, creating a form of hidden correlation that firm-level risk models, by construction, cannot detect and therefore cannot price.

**Insurance markets cannot diversify the pool.** With three providers controlling 63 per cent of the market, a single-provider outage can trigger simultaneous claims across a large fraction of the insured population, and reinsurers, recognising this aggregation exposure, respond with exclusions or sub-limits that erode effective coverage. The 90 per cent protection gap we documented in Section 2 is not a transitional market inefficiency awaiting correction but rather a structural feature of the risk itself. Resilience against correlated infrastructure failure is, in the economic sense, a public good, and markets will therefore systematically underprovide it.

**Government oversight assumes capacities dependency has already eroded.** Regulatory frameworks from the EU AI Act to US federal procurement guidelines require transparency, auditing, and compliance monitoring, each of which presupposes that the regulator possesses the technical capacity and jurisdictional authority to inspect the systems it oversees. Under conditions of infrastructural dependency, this presupposition fails because agencies cannot meaningfully audit systems whose source code they cannot access, running on infrastructure they cannot enter, maintained by personnel they do not employ. The state continues to perform oversight it can no longer meaningfully exercise, and this gap between procedural form and structural reality is itself a risk, because it sustains the appearance of democratic control while the prerequisites for such control have been transferred elsewhere.

Addressing these failures requires coordinated intervention on three fronts.

First, risk models for systemically important services should incorporate cloud concentration as a correlated risk factor, analogous to the treatment of geographic concentration in natural catastrophe modelling. Regulators of financial services and critical national infrastructure already mandate stress testing against a range of adverse scenarios, and there is no principled reason to exclude the failure of a dominant cloud provider from that set.

Second, a new insurance architecture is needed to bridge the protection gap. Parametric covers triggered by verified outage duration, structured along the lines of catastrophe bonds, could deliver rapid payouts without the claims-adjustment delays that currently impede conventional cyber insurance. For the systemic tail of the distribution, sovereign-backed facilities comparable to Pool Re for terrorism risk or Flood Re for flood risk may prove necessary, given the structural limits of private diversification in a market that is this concentrated.

Third, governments should pursue a strategy of bounded sovereignty, meaning targeted investment in sovereign cloud capacity for legitimacy-critical functions, combined with mandatory multi-cloud requirements and portability standards for systemically important public services. This is not an argument for infrastructural autarky, which would sacrifice genuine efficiency gains and is in any case impractical for most states. It is, rather, a risk management strategy that acknowledges the impossibility of full infrastructure independence while drawing a deliberate boundary around the functions where democratic legitimacy is most at stake and where the consequences of provider

failure extend beyond the economic into the constitutional.

## 6. Conclusion

Cloud infrastructure concentration has produced a sovereign-systemic risk that current risk management frameworks fail to capture, that insurance markets cannot efficiently transfer, and that government oversight mechanisms cannot effectively govern, and the exposure is compounding faster than the institutional capacity to address it.

The associated risk is of profound significance because path dependencies make remediation progressively more costly with each year of deeper integration. Switching costs rise, in-house technical capacity atrophies, and the set of viable alternatives narrows, which means that the cost of inaction is not constant but accelerating, exhibiting the hysteresis that makes delayed remediation progressively more expensive than early investment. The GDP-at-risk analysis we have presented suggests that the sovereign risk multiplier will double every three years through 2030, lending some urgency to the question of institutional design.

None of this constitutes a case against cloud computing, whose efficiency gains are genuine and whose global value has been estimated at $3 trillion in potential EBITDA by 2030 (McKinsey & Company, 2022). Governments are right to pursue those gains. Our argument is rather that states have digitised public authority faster than they have constitutionalised control over the infrastructure through which that authority is now exercised, and that this gap between digitalisation and governance is itself the risk. The Sovereign Infrastructure Exposure Index we propose here offers one instrument for closing that gap, a means by which any state can locate its position, identify the structural drivers of its exposure, and determine which interventions would most efficiently reduce it.

Whether that architecture is built deliberately or improvised under duress in the aftermath of a sovereign-scale disruption remains, for the moment, a political choice. But the risk itself is no longer hypothetical, and the window in which deliberate action remains affordable is, on the evidence we have assembled, closing.

## References


Ananny, M. and Crawford, K. (2018) 'Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability', *New Media & Society*, 20(3), pp. 973–989.

Gartner (2024) *Forecast: Public Cloud Services, Worldwide, 2022–2028.* Stamford, CT: Gartner, Inc.

Geneva Association (2024) *Cyber Insurance: Strengthening Resilience in a Digital Economy.* Zurich: The Geneva Association.

Guy Carpenter (2025) *Global Cyber Market Update: Q1 2025.* New York: Guy Carpenter & Company.

International Monetary Fund (2025) *World Economic Outlook: April 2025.* Washington, DC: IMF.

Mann, M. (1984) 'The autonomous power of the state: Its origins, mechanisms and results', *European Journal of Sociology*, 25(2), pp. 185–213.

McKinsey & Company (2022) 'Projecting the global value of cloud: $3 trillion is up for grabs for companies that go beyond adoption', *McKinsey Digital.*

McKinsey & Company (2024) 'Building cyber resilience in national critical infrastructure', *McKinsey Risk & Resilience.*

Marsh McLennan (2025) *Global Insurance Market Index: Q1 2025.* New York: Marsh McLennan.

OECD (2024) *Survey on the Implementation of the OECD AI Principles.* Paris: OECD Publishing.

Parametrix (2024a) *Cloud Downtime Index: Annual Report 2023.* New York: Parametrix Insurance.

Parametrix (2024b) *CrowdStrike-Windows Outage: Economic Impact Assessment*. New York: Parametrix Insurance.

Parametrix (2025a) *Cloud Downtime Index: Annual Report 2024*. New York: Parametrix Insurance.

Parametrix (2025b) *AWS US-East-1 Outage October 2025: Preliminary Impact Assessment*. New York: Parametrix Insurance.

Parametrix (2025c) *Azure Front Door Outage October 2025: Comparative Analysis*. New York: Parametrix Insurance.

Plantin, J.-C. et al. (2018) 'Infrastructure studies meet platform studies in the age of Google and Facebook', *New Media & Society*, 20(1), pp. 293–310.

Raji, I.D. et al. (2020) 'Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing', *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 33–44.

Star, S.L. and Ruhleder, K. (1996) 'Steps toward an ecology of infrastructure: Design and access for large information spaces', *Information Systems Research*, 7(1), pp. 111–134.

Synergy Research Group (2025) *Cloud Infrastructure Market Statistics: Q1 2025*. Reno, NV: Synergy Research Group.

Uptime Institute (2023) *Annual Outage Analysis 2023*. New York: Uptime Institute.