

The Materiality Gate: Dynamic Updating of AI Sovereignty Risk under Geopolitical Shocks

Subramanyam Sahoo

Independent

sahoo2vec@gmail.com

Abstract

Geopolitical risk is often scored too quickly. A contested presidential claim, a high profile military meeting, or a renamed command can alter expectations without changing who can acquire accelerators, operate a cloud region, migrate a model, or keep safety monitoring online. We introduce the *Materiality Gate*, a dynamic rubric that separates political signals from infrastructure shocks in national AI sovereignty assessments. An event changes an exposure score only when it is verified, exercises authority over a specific dependency, and has a documented path to change availability, substitutability, jurisdictional control, critical workload exposure, or safety continuity. We operationalize the gate across accelerators, semiconductor supply, cloud infrastructure, models and software, and physical infrastructure. We apply it to India during 2025 and 2026. The contested U.S. mediation claim after the May 2025 India–Pakistan conflict, the White House meeting with Pakistan’s army chief, Pakistan’s role in the Islamabad Memorandum of Understanding, and the restoration of the U.S. Pacific Command name all raise monitoring priority, but do not by themselves alter infrastructure scores. In contrast, the 2025 U.S. AI Diffusion Rule enters a pending material state because it created a legally specified access regime; its rescission before compliance shows why scoring must be versioned. The resulting method is conservative about rhetoric, responsive to real policy change, and symmetric: announcements are not assets, and symbols are not outages.

1 Introduction

AI sovereignty is increasingly discussed as a national property. A state is described as sovereign when it hosts data centers, funds domestic models, or announces access to large pools of accelerators. Yet modern AI systems are built from dependencies that cross territorial and corporate boundaries. A data center may sit inside a country while its accelerators, firmware, cloud control plane, model license, and replacement parts remain subject to actors elsewhere. Recent work therefore treats sovereignty as layered control rather than territorial possession alone [Lehdonvirta et al., 2024, Hawkins et al., 2025]. This view is consistent with research on weaponized interdependence, which shows how actors can exercise power through concentrated network positions [Farrell and Newman, 2019, Beaumier and Cartwright, 2024].

A second problem has received less attention: *when should a sovereignty score change?* Geopolitical narratives update every day. Infrastructure changes more slowly and through identifiable mechanisms. Without an update rule, analysts can convert political interpretation into technical risk. A diplomatic meeting may be treated as evidence that a country has lost access to compute. A government announcement may be treated as evidence that a dependency has disappeared. Both moves confuse a signal about future relations with a change in present capability.

The distinction becomes concrete in India. When a president says he stopped a war, a neighboring army chief is welcomed at the White House, and a military command drops “Indo” from its name, has India’s AI infrastructure become less secure? The answer cannot be inferred from the mood of the news cycle. It depends on whether a provider, regulator, contract, or supply chokepoint has actually changed the terms of access.

India is a useful diagnostic case. It is expanding public access to AI compute and building governance institutions, while remaining embedded in foreign accelerator, cloud, software, and semiconductor networks [Press Information Bureau, Government of India, 2026a,c]. It also sits inside a dense geopolitical environment involving Pakistan, China, the United States, the Indian Ocean, and long standing security disputes [Marshall, 2016]. Between May 2025 and June 2026, several events invited broad claims about a U.S. tilt toward Pakistan or a reduced strategic role for India. These included President Donald Trump’s disputed claim that he stopped the May 2025 India–Pakistan conflict, his meeting with Pakistan’s army chief, Pakistan’s mediation of an interim U.S.–Iran agreement, and the removal of “Indo” from the U.S. military command name [Ministry of External Affairs, Government of India, 2025, Reuters, 2025, 2026, U.S. Pacific Command, 2026]. Each event matters politically. None, on the public evidence available at the time of writing, directly changed access to an accelerator, a cloud region, a model, a semiconductor input, or a safety control.

This paper makes one claim and develops it in depth:

A geopolitical event should change an AI sovereignty score only when it passes a materiality gate linking verified action to a specific infrastructure dependency.

The gate does not ask whether a president meant what he said. It asks whether a researcher can still acquire an accelerator, move a workload, inspect an audit log, and keep a safety system online. This distinction is central to AI safety because compute and cloud providers can be both safety infrastructure and policy intermediaries [Sastry et al., 2024, Heim et al., 2024]. A false positive can provoke wasteful autarky or a race to build under lower safety standards. A false negative can leave critical evaluation and incident response dependent on an actor who can withdraw access.

We contribute a formal update rule, a five layer exposure rubric, and a worked India event sequence. The method assigns events to four states: WATCH, PENDING, ACTIVE, and WITHDRAWN. A score is recomputed only in the active state. The pending state records a promulgated or contracted change that has not yet taken operational effect. This is essential for policies that are announced, delayed, rescinded, or only partially implemented.

2 From chokepoints to an event update problem

2.1 AI sovereignty is dependency control

Compute governance research identifies several properties that make advanced computing a policy lever: it is comparatively measurable, excludable, and concentrated [Sastry et al., 2024]. Cloud providers can also record, verify, and enforce rules over customer workloads [Heim et al., 2024]. These properties create governance opportunities, but they also create dependence. Control may sit with the state where a data center is located, the company operating it, the jurisdiction governing that company, or the vendor producing the accelerator [Hawkins et al., 2025]. Digital platforms can acquire infrastructural power that resembles functions conventionally associated with states [Kelton et al., 2022].

We represent national AI infrastructure as a set of dependencies $d \in \mathcal{D}$. Each dependency belongs to one of five layers $\ell \in \mathcal{L}$:

1. compute accelerators;
2. the semiconductor stack, including design tools, fabrication, memory, equipment, and packaging;
3. cloud infrastructure;

4. models and software;
5. physical infrastructure, including power, cooling, facilities, and connectivity.

For every dependency and time t , we define an exposure vector

$$\mathbf{x}_{d,t} = (C_{d,t}, J_{d,t}, T_{d,t}, K_{d,t}, S_{d,t}) \in \{0, 1, 2, 3, 4\}^5, \quad (1)$$

where C is concentration, J is jurisdictional or external control, T is substitution or recovery time, K is critical workload exposure, and S is the consequence for safety continuity. Appendix A gives observable anchors for each layer. Equal weighting is the default:

$$E_{d,t} = \frac{1}{5} (C_{d,t} + J_{d,t} + T_{d,t} + K_{d,t} + S_{d,t}). \quad (2)$$

Other weights are permissible, but they must be declared before an event is coded. This prevents an analyst from changing weights to obtain a preferred geopolitical conclusion.

2.2 Why ordinary event coding fails

The usual response to uncertainty is to add a “geopolitical risk” variable. That move creates three problems. First, it bundles distinct mechanisms. An export rule, a presidential statement, a military clash, and a cloud outage are not commensurate events. Second, it rewards salience. The most visible event receives the largest update, even when the relevant dependency is unchanged. Third, it is asymmetric. Analysts often raise risk after threatening news but lower it immediately after investment announcements, even though announced capacity may not yet exist.

The result is narrative volatility. Scores become measures of attention rather than exposure. A useful rubric therefore needs two components: a stable set of dependency dimensions and a strict rule governing when events update them.

3 The Materiality Gate

3.1 Three tests

Let e be an event observed at time t . The Materiality Gate for dependency d is

$$\mathcal{G}(e, d) = V_e A_{e,d} M_{e,d}, \quad (3)$$

where each term is binary.

Verification, V_e . The event is established by a primary record, a legal instrument, an official operational notice, a contract, audited data, or multiple credible independent sources. Verification applies to the event itself, not automatically to every causal claim made about it. For example, a speech can be verified while the speaker’s account of mediation remains contested.

Authority exercised, $A_{e,d}$. The event exercises regulatory, ownership, contractual, operational, or coercive authority over dependency d . General political influence is insufficient. A head of government may possess broad authority, but a ceremonial meeting does not exercise export control authority over a chip order.

Material pathway, $M_{e,d}$. The evidence specifies how the event changes at least one component of Eq. 1. Examples include a license requirement, service suspension, provider withdrawal, verified activation of substitute capacity, tested failover, contract termination, loss of a network route, or an outage that degrades monitoring.

Figure 1 shows the update architecture. A failed test does not make an event irrelevant. It routes the event to a watchlist, records a review trigger, and may raise the priority of evidence collection. It does not change the infrastructure score.




Figure 1: The Materiality Gate. Political signals remain visible as review triggers, but scores change only after verified action exercises authority over a specific dependency through a documented material pathway.

3.2 Four event states

Passing the gate does not always imply an immediate score change. We assign each event to one of four states.

Watch. At least one gate condition fails. The event is logged with an evidence request and a next review date.

Pending. The gate passes, but the material consequence has a future effective date or depends on implementation. A promulgated export rule with a later compliance date is pending. A funded accelerator purchase is pending until delivery, commissioning, and user access are verified.

Active. The consequence is in force or operationally verified. The affected dimensions are recomputed from current evidence.

Withdrawn. A pending or active event is rescinded, expires, is enjoined, or is operationally reversed. The score returns only after the reversal is verified. The historical record is preserved.

This state machine matters because the time between announcement and operation is often where policy changes. It also prevents a positive announcement from being counted twice, first when announced and again when deployed.

3.3 Evidence and confidence

Every coded observation contains an as of date, source, rationale, confidence grade, trend, and review trigger. Confidence is reported separately from exposure. Grade A denotes an official, contractual, audited, or directly verified record. Grade B requires at least two current credible sources. Grade C denotes one credible but incomplete or indirect source. Grade D denotes an estimate or expert inference. Missing evidence is marked “not assessed,” never scored as zero.

The separation between score and confidence is important. A dependency may be critically exposed with low confidence, or resilient with high confidence. Folding confidence into the score would make the meaning of the number ambiguous.

4 India as a diagnostic case, 2025–2026

4.1 Case selection and baseline

India combines growing domestic capability with deep interdependence. Official statements in February 2026 reported more than 38,000 high end GPUs onboarded through the IndiaAI Mission and announced 20,000 additional GPUs [Press Information Bureau, Government of India, 2026c]. A March statement described the additional 20,000 as “currently under process” [Press Information Bureau, Government of India, 2026b]. The Materiality Gate records the first quantity according to the scope and definition in the official source, while assigning the additional quantity to PENDING. It does not reduce exposure until delivery, commissioning, allocation terms, and actual availability are verified.

This domestic example establishes the symmetry of the method. Favorable announcements face the same test as adverse geopolitical signals. The question is not whether a policy sounds reassuring. The question is whether a dependency has changed.

The case does not attempt to produce a definitive national score. Public data are insufficient for that task, especially for workload shares, contract terms, replacement inventories, firmware control, and tested failover. Instead, we use the event sequence to test whether a transparent analyst would update the rubric and why. The evidence freeze is 21 June 2026. Later developments should enter as new event records rather than being silently backfilled into earlier scores.

4.2 Event sequence

Figure 2 places the events on a common timeline. The January 2025 AI Diffusion Rule is the positive control. It specified a worldwide licensing framework for advanced computing integrated circuits and certain model weights [U.S. Department of Commerce, Bureau of Industry and Security, 2025b]. It therefore passes the gate for accelerator and model dependencies. Because compliance requirements were scheduled for May 15, it first entered PENDING. The Commerce Department announced on May 13 that it would not enforce the new requirements and would rescind the rule [U.S. Department of Commerce, Bureau of Industry and Security, 2025a]. The pending shock was therefore moved to WITHDRAWN before its compliance date.

The remaining events are politically salient but fail at least one dependency test. After the May 2025 India–Pakistan conflict, Trump repeatedly claimed that he had stopped the fighting, including through trade pressure. India stated that the cessation was arranged through existing military channels and that no U.S. mediation or trade discussion occurred [Ministry of External Affairs, Government of India, 2025, Reuters, 2025]. The claim is therefore contested. More importantly for this paper, neither account documents a change in an AI infrastructure dependency.

On June 18, 2025, Trump hosted Pakistan’s army chief, Field Marshal Asim Munir, at the White House [Reuters, 2025]. The meeting can reasonably raise the monitoring priority assigned to U.S.–Pakistan relations. It does not establish preferential accelerator access, a cloud service change, an export license denial to India, or a reduction in safety cooperation. It remains in WATCH.

In June 2026, Pakistan acted as witness and mediator for the interim U.S.–Iran agreement titled the Islamabad Memorandum of Understanding [Reuters, 2026]. This elevated Pakistan’s diplomatic role. It again lacked a documented transmission path to India’s AI dependencies and remains in WATCH.

Finally, the U.S. military restored the U.S. Pacific Command name in June 2026. The official command biography reflected the new designation by June 18 [U.S. Pacific Command, 2026]. The 2018 addition of “Indo” had explicitly recognized the growing connection between the Indian and Pacific Oceans [U.S. Department of Defense, 2018]. Reversing that label is a strategic signal, but a label is not force posture, intelligence sharing, an export rule, or a compute contract. The event remains in WATCH unless operational evidence follows.




Figure 2: Event states in the India case. Only the export rule and the capacity announcement enter a pending material state. The rule is withdrawn before compliance; the announced capacity remains pending until operational evidence is available.

Event	V	A	M	Rubric consequence
AI Diffusion Rule issued	1	1	1	PENDING for accelerator and model access; recompute at effect or verified implementation
Rule rescission announced	1	1	1	Move pending shock to WITHDRAWN; retain event history
Trump mediation and trade claim	1	0	0	WATCH; record dispute, seek dependency evidence, no score change
Munir White House meeting	1	0	0	WATCH; raise diplomatic monitoring priority only
20,000 additional GPUs announced	1	1	1	PENDING capacity; do not reduce exposure before operation and access are verified
Islamabad MOU	1	0	0	WATCH; diplomatic elevation without documented AI infrastructure effect
Pacific Command name restored	1	0	0	WATCH; symbolic signal unless posture or technology policy changes

Table 1: Application of the Materiality Gate. *V* is event verification, *A* is dependency specific authority exercised, and *M* is a documented material pathway.

4.3 Gate results

Table 1 reports the coding. “Authority exercised” is dependency specific. The U.S. president has broad political authority, but a speech or meeting does not itself exercise export, contractual, or operational control over an accelerator order. This distinction prevents the actor’s office from substituting for a causal mechanism.

The positive control demonstrates that the gate is not a device for dismissing geopolitics. The AI Diffusion Rule would have changed the legal conditions governing advanced computing access. It receives stronger treatment than every diplomatic event in the sequence. The rescission then demonstrates why the system must preserve state transitions rather than overwrite the past.

The domestic capacity announcement provides a second control. It tests favorable evidence. A planned expansion can be highly credible and still not be operational. Counting it immediately would understate concentration and substitution risk. The gate therefore resists both pessimistic and optimistic narrative pressure.

5 Implications for AI safety and middle power strategy

5.1 Safety continuity is a first class exposure dimension

Traditional sovereignty indexes focus on ownership, location, or aggregate capacity. AI safety requires an additional question: what happens to evaluation, monitoring, incident response, and control during disruption? This motivates the *S* dimension in Eq. 1. The relevant unit is not simply compute quantity. It is the continuity of safety relevant functions.

Consider two accelerator shortages with equal capacity loss. In the first, commercial training slows while red team evaluations and incident response retain reserved capacity. In the second, all workloads compete for the same pool and safety evaluations are deferred. The geopolitical event is similar, but the safety exposure differs. Likewise, a cloud migration that preserves data but loses audit logs, model provenance, or emergency credentials is not a safe failover.

This framing follows the socio technical view that safety failures emerge through institutions and deployment conditions, not only model behavior [Amodei et al., 2016, Raji and Dobbe, 2023]. It also makes the rubric actionable. Governments can reduce exposure through protected evaluation capacity, tested portability, audit log escrow, key management separation, and contract clauses preserving access for incident response.

5.2 Resilient interdependence, not reflexive autarky

A score that reacts to every diplomatic signal can make autarky appear rational. Yet complete self sufficiency across accelerators, fabrication, memory, design tools, clouds, models, energy, and cables is infeasible for most states and may produce new concentration. A rushed sovereignty program can also redirect scarce compute from independent safety research toward prestige models.

The Materiality Gate supports a narrower strategy: identify dependencies where authority and material pathways are real, then reduce single points of failure. This may involve supplier diversity, multi jurisdiction contracts, tested workload portability, domestic maintenance capability, strategic inventories, and reciprocal emergency access. The goal is resilient interdependence, not the absence of foreign inputs.

For middle powers, this is a practical distinction. Diplomatic alignment can change quickly. Infrastructure contracts and supply chains change through slower, observable acts. A versioned dependency register allows a state to cooperate widely while preparing for specific chokepoints.

5.3 A dashboard should explain why it did not move

Most monitoring systems emphasize changes. A trustworthy sovereignty dashboard should also explain nonchanges. When a salient event fails the gate, the system should publish the failed condition, the evidence sought, and the review trigger. For the Pacific Command renaming, for example, triggers might include changes in Indian Ocean force posture, Quad exercises, defense technology agreements, export treatment, or intelligence arrangements. For Pakistan's diplomatic elevation, triggers might include preferential access to controlled technology that materially changes India's substitution environment.

This approach does not claim that symbols are meaningless. Symbols can precede policy. It treats them as hypotheses requiring a transmission path. That makes the analysis responsive without allowing interpretation to masquerade as measurement.

6 Limitations

The method has five limitations. First, public evidence is incomplete. Contractual controls, provider dependencies, spare inventories, and cyber access may not be observable. The confidence field makes this limitation visible but cannot remove it.

Second, the gate is conservative. It may delay an update when credible intelligence indicates an impending restriction before a public act. Analysts with classified or contractual evidence can apply the same method, but their result may not be reproducible publicly.

Third, the 0 to 4 anchors discretize continuous risk. They are intended for transparent comparison and sensitivity analysis, not precise probability estimation. The underlying observations should be retained so later work can fit probabilistic models.

Fourth, dependencies interact. A semiconductor bottleneck can cascade into cloud availabil-

ity and model substitution. Appendix A includes a bottleneck rule for the semiconductor layer, but a complete dependency graph is future work.

Fifth, this paper does not adjudicate broad claims about whether U.S. policy “favors” India or Pakistan. Favor is not an infrastructure variable. The method instead specifies which observable acts would make such a shift relevant to AI sovereignty.

7 Conclusion

Geopolitical analysis becomes useful for AI safety when it reaches a dependency. The Materiality Gate provides a simple discipline: verify the event, identify the authority exercised, and document the material pathway. Events that fail remain visible as watchlist signals. Events with future consequences enter a pending state. Only active, verified consequences change exposure scores.

The India case shows why this discipline matters. A disputed mediation claim, a high profile meeting, an international accord named for Islamabad, and a command renaming can reshape political expectations without changing present AI infrastructure. A licensing rule can change infrastructure access, even before a single chip shipment is denied. A GPU announcement can improve future resilience without improving current resilience. Treating these events differently is not caution for its own sake. It is measurement.

The paper’s central rule is intentionally plain: announcements are not assets, and symbols are not outages. A winning sovereignty strategy should be equally plain about what has actually changed.

Ethics statement

This work analyzes public geopolitical events that are politically contested. We separate verified events from disputed causal claims and avoid inferring national intent from rhetoric alone. The rubric is designed to reduce partisan overinterpretation by requiring a dependency specific material pathway before a score changes. It uses public, aggregate information and does not identify sensitive facility vulnerabilities, operational security details, or exploitable failure points. The method could still be misused to present subjective weights as objective rankings. We therefore require declared weights, source provenance, confidence labels, and a preserved change log.

References

- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in AI safety. *arXiv preprint arXiv:1606.06565*, 2016. doi: 10.48550/arXiv.1606.06565.
- Guillaume Beaumier and Madison Cartwright. Cross-network weaponization in the semiconductor supply chain. *International Studies Quarterly*, 68(1):sqae003, 2024. doi: 10.1093/isq/sqae003.
- Henry Farrell and Abraham L. Newman. Weaponized interdependence: How global economic networks shape state coercion. *International Security*, 44(1):42–79, 2019. doi: 10.1162/isec.a.00351.
- Zoe Jay Hawkins, Vili Lehdonvirta, and Boxi Wu. AI compute sovereignty: Infrastructure control across territories, cloud providers, and accelerators. Technical report, Social Science Research Network, June 2025.
- Lennart Heim, Tim Fist, Janet Egan, Sihao Huang, Stephen Zekany, Robert Trager, Michael A. Osborne, and Noa Zilberman. Governing through the cloud: The intermediary role of compute providers in AI regulation. *arXiv preprint arXiv:2403.08501*, 2024. doi: 10.48550/arXiv.2403.08501.

- Maryanne Kelton, Michael Sullivan, Zac Rogers, Emily Bienvenue, and Sian Troath. Virtual sovereignty? private internet capital, digital platforms and infrastructural power in the united states. *International Affairs*, 98(6):1977–1999, 2022. doi: 10.1093/ia/iia226.
- Vili Lehdonvirta, Bóxi Wú, and Zoe Hawkins. Compute north vs. compute south: The uneven possibilities of compute-based AI governance around the globe. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, pages 828–838, 2024. doi: 10.1609/aies.v7i1.31683.
- Tim Marshall. *Prisoners of Geography: Ten Maps That Tell You Everything You Need to Know About Global Politics*. Elliott and Thompson, London, 2016. ISBN 9781783962433.
- Ministry of External Affairs, Government of India. English translation of foreign secretary’s statement on the telephone conversation between the prime minister and the U.S. president, June 2025. URL https://www.mea.gov.in/Speeches-Statements.htm?dtl/39692/English_Translation_of_Foreign_Secretarys_statement_on_the_telephone_conversation_between_PM_and_US_President_June_17_2025.
- Press Information Bureau, Government of India. India AI governance guidelines, 2026a. URL <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2228315>.
- Press Information Bureau, Government of India. AI models developed under IndiaAI mission represent significant step towards technological self-reliance, March 2026b. URL <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2239614>.
- Press Information Bureau, Government of India. India’s AI moment, February 2026c. URL <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2229397>.
- Inioluwa Deborah Raji and Roel Dobbe. Concrete problems in AI safety, revisited. *arXiv preprint arXiv:2401.10899*, 2023. doi: 10.48550/arXiv.2401.10899.
- Reuters. Trump hosts pakistani army chief, disagrees with india over india-pakistan war mediation, June 2025. URL <https://www.reuters.com/world/asia-pacific/india-will-not-accept-third-party-mediation-relations-with-pakistan-modi-tells-2025-06-18/>.
- Reuters. White house sends text of interim U.S.-Iran agreement to U.S. congress, June 2026. URL <https://www.reuters.com/world/middle-east/white-house-sends-text-interim-us-iran-agreement-us-congress-2026-06-18/>.
- Girish Sastry, Lennart Heim, Haydn Belfield, Markus Anderljung, Miles Brundage, Julian Hazell, Cullen O’Keefe, Gillian K. Hadfield, Richard Ngo, Konstantin Pilz, George Gor, Emma Bluemke, Sarah Shoker, Janet Egan, Robert F. Trager, Shahar Avin, Adrian Weller, Yoshua Bengio, and Diane Coyle. Computing power and the governance of artificial intelligence. *arXiv preprint arXiv:2402.08797*, 2024. doi: 10.48550/arXiv.2402.08797.
- U.S. Department of Commerce, Bureau of Industry and Security. Department of commerce announces rescission of biden-era artificial intelligence diffusion rule, strengthens chip-related export controls, May 2025a. URL <https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens>.
- U.S. Department of Commerce, Bureau of Industry and Security. Framework for artificial intelligence diffusion, January 2025b. URL <https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion>. 90 Federal Register 4544, published 15 January 2025.
- U.S. Department of Defense. Pacific command change highlights growing importance of indian ocean area, May 2018. URL <https://www.defense.gov/News-Stories/Article/Article/1535808/pacific-command-change-highlights-growing-importance-of-indian-ocean-area/>.
- U.S. Pacific Command. Commander, U.S. pacific command, June 2026. URL <https://www.pacom.mil/Contact/Directory/Bio-Display/Article/2590636/commander-us-indo-pacific-command/>. Official command biography, last updated 18 June 2026.

A Full dynamic exposure rubric

All criteria use scores from 0 to 4. Tables give anchors for 0, 2, and 4. Scores 1 and 3 represent evidence between adjacent anchors. Each score must be accompanied by an as of date, evidence, rationale, confidence grade, trend, and review trigger.

A.1 Compute accelerators

Criterion	0: resilient	2: moderate	4: critical
Supplier concentration	Largest vendor or supplier jurisdiction provides less than 25% of usable capacity; at least three viable platforms	Largest dependency provides 40–60%; two practical platforms	More than 80% comes from one vendor or jurisdiction, or one platform is indispensable
Foreign control	No single foreign actor can materially restrict use, repair, updates, support, or procurement	Some capacity is subject to foreign licensing or provider discretion, but operational alternatives exist	One foreign authority or provider can substantially restrict access, use, maintenance, firmware, or procurement
Substitution time	Equivalent capacity can be activated within one month	Substitution requires three to six months	Substitution requires more than 18 months, or no equivalent exists
Critical workload exposure	Disruption affects less than 10% of nationally important AI workloads	Disruption affects 25–50%	Disruption affects more than 75%, or disables an essential function
Safety continuity	Mainly a cost or performance effect; evaluation and monitoring continue	Safety evaluation, red teaming, or research is materially delayed	Monitoring, incident response, model control, or critical safety functions become unavailable

Table 2: Rubric for compute accelerators.

A.2 Semiconductor stack

The stack includes design tools, semiconductor intellectual property, materials, fabrication, production equipment, advanced packaging, and high bandwidth memory.

Criterion	0: resilient	2: moderate	4: critical
Chokepoint concentration	Every essential stage has at least three qualified sources across jurisdictions	One or two stages have only two practical sources	An indispensable stage has one qualified source or is concentrated in one jurisdiction
Jurisdictional exposure	No single foreign jurisdiction can block the complete supply path	Several inputs require foreign approval, but alternate paths remain	One jurisdiction or aligned group can halt the complete supply path
Requalification time	Alternative supplier or process can be qualified within three months	Qualification requires six to twelve months	Qualification requires more than two years, or no substitute exists
Supply impact	Disruption removes less than 10% of relevant supply	Disruption removes 25–50%	Disruption removes more than 75% or eliminates access to the required technology class
Safety and trust	Alternatives preserve reliability, provenance, cyber security, and hardware assurance	Alternatives reduce assurance or traceability	Substitution requires unaudited, counterfeit prone, insecure, or unreliable hardware, or prevents safety critical computing

Table 3: Rubric for the semiconductor stack. If any indispensable stage receives a 4, the overall layer score cannot be below 3.

A.3 Cloud infrastructure

Criterion	0: resilient	2: moderate	4: critical
Provider concentration	Largest provider hosts less than 25% of relevant workloads; tested multi provider deployment exists	Largest provider hosts 40–60%; partial redundancy exists	More than 80% depends on one provider, or one provider performs an irreplaceable function
Legal and contractual control	Data, keys, and control functions are locally controlled or distributed; suspension rights are constrained	Foreign law or provider terms affect some services, but fallback exists	A provider or foreign authority can unilaterally suspend service, restrict access, or control critical data or systems
Portability and failover	Tested failover occurs within 24 hours	Migration requires one to four weeks and material reconfiguration	Migration requires more than three months, is contractually obstructed, or is technically infeasible
Critical workload exposure	Less than 10% of critical workloads are interrupted	25–50% are interrupted	More than 75%, or an essential public service, is interrupted
Safety continuity	Audit logs, backups, identity controls, evaluations, and emergency controls survive failover	Some telemetry, provenance, or monitoring is temporarily lost	Failover causes loss of audit records, access controls, safety monitoring, provenance, or emergency shutdown capability

Table 4: Rubric for cloud infrastructure.

A.4 Models and software

Criterion	0: resilient	2: moderate	4: critical
Provider and license concentration	Critical functions run on at least three interoperable models or software stacks	Two alternatives exist, but switching involves substantial lock in	One API, model, license, or framework is indispensable
External access control	Models can be locally deployed, use open standards, or have durable access rights	Provider policy, model updates, or license conditions materially affect operation	An external provider can terminate access, change behavior, revoke rights, or withhold essential safety information
Substitution time	Replacement and validation require less than one week	Replacement, retraining, or validation requires one to three months	Replacement requires more than six months, or no functional equivalent exists
Critical task exposure	Less than 10% of critical workflows depend on the stack	25–50% depend on it	More than 75%, or an essential public function, depends on it
Safety equivalence	Alternatives pass equivalent safety evaluation and remain auditable	Replacement introduces uncertain regressions or reduced transparency	Available replacements are materially less safe, cannot be audited, or cannot reproduce necessary safeguards

Table 5: Rubric for models and software.

A.5 Physical infrastructure

Criterion	0: resilient	2: moderate	4: critical
Geographic and asset concentration	Capacity is distributed across regions, grids, facilities, and routes; no asset supports more than 25%	40–60% is concentrated in one region, grid, facility cluster, or route	More than 80% depends on one site, grid zone, cable route, cooling system, or other chokepoint
Ownership and control	Essential assets and remote control systems are locally controlled or diversified	Some essential assets, spares, or management systems are foreign controlled	One foreign actor controls an indispensable asset, spare supply, remote management system, or network route
Recovery time	Automatic failover occurs within 24 hours and full repair within one month	Failover requires one to four weeks or repair one to six months	Recovery requires more than six months, or no practical replacement exists
Capacity affected	Failure removes less than 10% of relevant capacity	Failure removes 25–50%	Failure removes more than 75% or disables an essential national service
Safety impact	Systems shut down safely and monitoring remains functional	Monitoring, backups, or emergency operations are partly degraded	Failure causes uncontrolled shutdown, data loss, loss of monitoring, unsafe restoration, or interruption of safety critical systems

Table 6: Rubric for physical infrastructure.

B Event coding protocol

For each event and dependency pair, analysts should complete the following record before changing a score:

Field	Required content
Event identifier	Stable name, date, and jurisdiction
Dependency object	Specific provider, product, facility class, contract, supply stage, or safety function
Verification test	Primary source or two independent sources establishing the event
Authority test	Legal, contractual, ownership, operational, or coercive authority exercised over the dependency
Material pathway	Dimension changed, causal link, effective date, and measurable consequence
State	Watch, pending, active, or withdrawn
Prior/revised values	C, J, T, K, S before and after the event; blank if no update
Confidence	A, B, C, or D, reported separately from exposure
Review trigger	Concrete evidence that would reopen the coding decision
Analyst note	Alternative explanation and known missing data

Table 7: Minimum event record for a reproducible update.

B.1 Decision rules

1. Score operational capacity separately from announced or contracted capacity.
2. Do not infer a material pathway from diplomatic tone alone.
3. Do not treat missing data as resilience.
4. Preserve prior scores and event states. Never overwrite the change history.
5. Reassess quarterly and immediately after export control changes, sanctions, conflict escalation, major outages, contract changes, acquisitions, verified commissioning, or supplier withdrawal.
6. Run an equal weight sensitivity analysis if policy weights are used.
7. For sequential supply chains, apply bottleneck overrides before averaging.

C Worked event records

C.1 Disputed mediation claim

Event. Trump stated that he stopped the May 2025 India–Pakistan conflict and linked the outcome to trade. India stated that the cessation was arranged directly through military channels and rejected third party mediation [[Ministry of External Affairs, Government of India, 2025](#), [Reuters, 2025](#)].

Gate. The public statement is verified. It does not exercise dependency specific authority and supplies no documented material pathway. State: WATCH. Score effect: none.

Review triggers. A change in export licensing for India or Pakistan; preferential access to controlled accelerators; suspension of a cloud or model service; a defense technology agreement that changes substitution time; or a verified reduction in safety cooperation.

C.2 AI Diffusion Rule

Event. The January 2025 rule created a new framework governing advanced computing integrated circuits and certain model weights, with later compliance requirements [[U.S. Department of Commerce, Bureau of Industry and Security, 2025b](#)].

Gate. Verification, authority, and material pathway all pass. State at issuance: PENDING. Affected dimensions: accelerator and model jurisdictional exposure, availability, and substitution time.

State transition. BIS announced rescission and nonenforcement before the compliance date [U.S. Department of Commerce, Bureau of Industry and Security, 2025a]. State: WITHDRAWN. The pending update is canceled, but the event remains in the history because it reveals a plausible policy channel.

C.3 Additional IndiaAI capacity

Event. The government announced 20,000 additional GPUs in February 2026; a March statement described the expansion as under process [Press Information Bureau, Government of India, 2026c,b].

Gate. The announcement is verified and linked to an authorized program. The material pathway is procurement and commissioning, but the consequence is not yet operationally established. State: PENDING. Score effect: none until verified operation and access.

Review triggers. Delivery records, commissioning, provider and accelerator composition, allocation rules, observed user access, and tested availability during disruption.

D Suggested data schema

A minimal machine readable implementation uses one row per dependency criterion and observation date:

```
country, layer, dependency_id, criterion,  
score_0_to_4, as_of_date, event_id, event_state,  
source, rationale, confidence, trend, review_trigger
```

A separate event table stores V , A , and M , allowing analysts to reproduce why an event did or did not change a score. This separation is important: the event log can update frequently while the infrastructure table remains stable.