

THE THREAT OF EMERGING PRIVATE STRATEGIC POWER

What to Watch For — and What to Do About It

A policy white paper on AI, robotics, defense technology, communications, capital,
and the rise of non-state strategic actors

Prepared September 2026
Analytical draft for public-policy discussion

Executive Summary

A new category of power is emerging. A small number of private technology networks increasingly combine capabilities that, until recently, were found mainly inside sovereign states: frontier artificial intelligence, global communications, space launch, intelligence software, battlefield systems, autonomous machines, vast pools of capital, and direct access to political leaders.

This paper calls that combination private strategic power. It does not allege that any named executive or company is pursuing world domination. It argues something more defensible and more important: when privately controlled capabilities become strategically decisive, democratic societies should not rely on the current owner's goodwill, competence, or permanence. Institutions must be designed for the worst plausible successor, not the most reassuring incumbent.

The principal risk is convergence. An AI company by itself is not a sovereign rival. Neither is a satellite operator, robot manufacturer, defense contractor, venture fund, or politically connected investor. But a coordinated network combining intelligence, physical agency, communications, military integration, energy, finance, and jurisdictional shelter could acquire bargaining power greater than that of many states—without a state's constitutional duties, public-law transparency, electoral accountability, or territorial responsibilities.

The appropriate response is neither blanket nationalization nor unrestricted technological decentralization. It is an AI-era separation of powers: distributed ownership and authority; controlled access to catastrophic capabilities; independent countervailing systems; mandatory interoperability and continuity safeguards; auditable government procurement; restrictions on cross-domain concentration; and credible public authority over systems essential to national survival.

Core finding	Immediate priority	Long-term objective
Concentration—not proven malign intent—is the primary structural vulnerability.	Map cross-domain control, dependencies, and single points of failure.	Ensure no private actor, state, or alliance holds an uncontested catastrophic capability.

1. The Concept: From Corporate Influence to Strategic Power

Large corporations have always influenced governments. Private strategic power is different in kind, not merely degree. It exists when a private actor can independently shape the security choices of states, deny or condition access to essential infrastructure, act at machine speed across borders, or combine digital intelligence with physical force.

Five characteristics distinguish it from ordinary market power:

- Cross-domain integration: control or influence spanning AI, data, communications, logistics, defense, finance, energy, biotechnology, or robotics.

- **Strategic indispensability:** governments or populations cannot quickly substitute another supplier without accepting serious security or economic harm.
- **Operational autonomy:** the actor can deploy capability without waiting for ordinary public authorization.
- **Opacity:** ownership, technical behavior, government relationships, or internal decision rules are difficult for outsiders to examine.
- **Jurisdictional arbitrage:** key assets can be distributed across friendly or permissive jurisdictions, complicating effective oversight.

No single factor is decisive. The danger arises when several mature together under common ownership, contractual coordination, aligned financing, or durable personal networks.

2. Why This Is Historically New

States retain enormous countervailing advantages: law, taxation, police, armed forces, borders, central banks, intelligence services, and the power to seize or regulate physical assets. Claims that a technology billionaire is already “more powerful than a state” therefore need qualification. Yet technology has changed what private actors can do before a state can respond—and how costly intervention may become.

Software scales globally; autonomous systems can act faster than bureaucratic review; satellite constellations cross borders; cloud and AI services may become embedded in military and civil administration; and private capital can build strategic infrastructure outside normal appropriations cycles. Dependency can arise gradually through procurement, convenience, and superior performance, long before lawmakers recognize that an ordinary vendor has become a constitutional problem.

The key transition is from supplier to gatekeeper. A supplier sells a replaceable product. A gatekeeper can determine whether a state sees, communicates, decides, moves, or fights.

3. The Capability Stack

Domain	Strategic function	Risk when concentrated	Illustrative safeguard
Frontier AI	Prediction, planning, code, research	Information dominance and accelerated coercive capability	Independent evaluations; secure access tiers
Data/intelligence	Population and battlefield awareness	Private visibility into state and citizen activity	Purpose limits; logs; public audit authority
Communications/space	Connectivity, sensing, navigation	Ability to condition national or military operations	Continuity rights; multi-provider redundancy

Robotics/autonomy	Physical action at scale	Digital intent becomes persistent physical agency	Licensing, remote-disable governance, weapons rules
Defense integration	Targeting and command support	Vendor becomes embedded in sovereign force decisions	Human command responsibility; modular procurement
Compute/chips/energy	Underlying capacity	Competitors and regulators lose practical leverage	Capacity disclosure; diversified supply; emergency access
Capital/network control	Fast coordination across firms	Nominally separate entities act as one strategic bloc	Beneficial-control mapping; cross-domain review
Biotechnology	Design and automated experimentation	Catastrophic biological capability	Synthesis screening; lab assurance; defensive surveillance

4. Case Study: A Network, Not a Conspiracy

The loose ecosystem popularly called the “PayPal Mafia” is a useful stress test because alumni and associates have founded, funded, or influenced firms across software, defense, space, AI, finance, and politics. The label can mislead: it does not establish unified command, shared intent, or operational coordination. Networks often contain rivalry and disagreement. The correct analytical question is not whether a secret committee exists, but whether aligned capital, personnel, ideology, contracts, and infrastructure can produce coordinated power even without one.

Palantir occupies an unusually strategic position because its software supports government, intelligence, and military workflows. SpaceX and Starlink demonstrate how privately operated infrastructure can become operationally significant to national defense; a 2025 U.S. approval contemplated \$150 million in Starlink-related satellite communications support for Ukraine. Tesla publicly presents autonomy and humanoid robotics as central to bringing AI into the physical world. Anduril and related defense-technology firms emphasize autonomous platforms and software-defined military systems.

These facts justify scrutiny of dependency and convergence. They do not prove a plan for domination. Indeed, present humanoid robots remain far from a mass autonomous army; factories, launch sites, employees, banks, chip supplies, and radio spectrum remain subject to sovereign control. A disciplined paper must preserve both sides of that ledger.

5. Argentina: Relevant Signal, Insufficient Proof

President Javier Milei received Peter Thiel, representatives of Thiel Capital, and a Founders Fund partner at Casa Rosada on 23 April 2026. That official meeting establishes high-level access. It does not establish special territorial rights, control of Argentine institutions, an undisclosed Palantir agreement, or a sovereign sanctuary.

Argentina nevertheless belongs on a watch list for a narrower reason. A friendly jurisdiction can offer land, energy, data infrastructure, investment protection, regulatory accommodation, and political access. Concern should rise only if ordinary investment evolves into sovereign-like privileges, security autonomy, communications insulation, special legal immunities, or coordinated relocation of strategic assets. The evidentiary threshold should be concrete instruments—laws, contracts, concessions, licenses, land grants, force arrangements—not atmosphere or ideological affinity.

6. How to Evaluate Claims Without Sliding Into Paranoia

Existential-risk analysis must examine low-probability, high-consequence scenarios. But it must also resist a common error: treating every ambiguous event as confirmation of a prior fear. The solution is an explicit evidence ladder.

Category	Standard	How it should be written
Verified fact	Primary document or multiple reliable sources	State directly, with source and date.
Supported inference	Several facts make an explanation more likely	Identify it expressly as an inference and name alternatives.
Allegation	Identifiable person or report asserts it	Attribute it; do not repeat it as fact.
Speculation	Possible but presently unsupported	Use only to define indicators that could later confirm or falsify it.

A serious assessment should also record disconfirming evidence: technical delays, failed scaling, government resistance, loss of contracts, antitrust action, leadership conflict, supply dependence, and public rejection. A warning system that can only move its probability upward is not analysis; it is belief protection.

7. Ten Indicators That Would Materially Increase Concern

1. Sovereign-like territorial privileges. Long-term zones, immunities, or governance rights that constrain ordinary national jurisdiction.

2. Private force or autonomous weapons at scale. Operational control, not merely manufacture, of armed autonomous systems or security forces.

- 3. Compute and semiconductor independence.** Vertically integrated access to advanced chips, fabrication, cloud capacity, and model training insulated from export controls.
- 4. Independent energy and industrial capacity.** Power generation and production sufficient to sustain strategic operations through state pressure.
- 5. Mass autonomous-robot production.** Demonstrated volume, reliability, independent navigation, coordinated control, and adaptable physical work—not aspirational factory capacity.
- 6. Insulated biotechnology.** High-containment or automated biological R&D outside credible independent oversight.
- 7. Intervention-resistant communications.** Networks architected to remain controllable by the private operator despite lawful national intervention.
- 8. Jurisdictional relocation.** Coordinated movement of leadership, compute, factories, financial assets, and research into unusually permissive states.
- 9. Sanctions-resistant finance.** Payment, custody, and ownership structures capable of surviving coordinated government restrictions.
- 10. Substitution for government.** Private provision of identity, adjudication, currency, intelligence, security, or essential administration coupled with political efforts to displace public authority.

The decisive signal is convergence. One or two indicators may describe an ambitious multinational. Seven or eight under coordinated control could describe an emerging strategic actor. Analysts should score capability, coordination, scale, insulation, and intent separately rather than collapse them into one dramatic judgment.

8. What Would Decrease Concern

- Genuine technical failure or persistent inability to scale robotics, autonomy, or frontier models into reliable operations.
- Open standards, portability, interoperability, and credible multi-provider substitution for strategic services.
- Independent boards, dispersed ownership, transparent related-party arrangements, and meaningful separation among network companies.
- Government continuity plans that can operate without the vendor and are exercised in practice.
- Binding acceptance of external audits, incident reporting, democratic oversight, and lawful emergency authority.
- Visible political pluralism within the network rather than coordinated capture across jurisdictions.

9. The Policy Response: An AI Separation of Powers

The aim should not be to punish scale or political views. It should be to prevent any actor—private or public—from combining intelligence, infrastructure, and coercive capacity without countervailing control.

9.1 Map strategic control, not merely corporate ownership

Competition and national-security reviews should examine common investors, directors, contractual exclusivity, data sharing, key-person authority, technical integration, and reciprocal dependencies. Nominally separate companies can form a single functional capability stack. Governments need a continuously updated strategic-dependency register covering AI, cloud, satellites, launch, defense software, autonomous systems, energy, and biotechnology.

9.2 Require redundancy and exit rights

No essential military, emergency, communications, or civil-administration function should depend on a single private platform without a tested substitute. Strategic contracts should include interoperability, data portability, escrow where appropriate, continuity of service, transition assistance, and clearly bounded emergency authorities. Redundancy costs money; dependence can cost sovereignty.

9.3 Separate catastrophic capabilities

A single private control center should not simultaneously command frontier intelligence, critical communications, mass physical autonomy, and high-consequence biological or military systems. Cross-domain concentration review should be triggered before acquisitions or technical integrations create that combination. Structural separation, access firewalls, independent governance, and licensing may be more effective than after-the-fact conduct rules.

9.4 License deployment, not ideas

Regulation should focus on the point where capability becomes strategically consequential: training or operating the most capable models; connecting autonomous systems to weapons or critical infrastructure; deploying intervention-resistant communications; or automating high-risk biological experimentation. This protects ordinary research and open debate while controlling dangerous scale and coupling.

9.5 Build independent countervailing AI

Democratic governments and alliances need technically independent systems able to audit, detect, attribute, and counteract catastrophic behavior by other systems. Diversity matters: systems built by different teams, trained differently, and supplied with different intelligence are less likely to share the same blind spots. The EU AI Act is an early step through documentation, systemic-risk evaluation, mitigation, incident reporting, and cybersecurity obligations for the most capable general-purpose models. It is not yet a complete separation-of-powers regime.

9.6 Preserve public command responsibility

AI may advise military, intelligence, public-health, and emergency decisions, but accountable public officials must remain legally responsible for their use. Governments should not allow

vendors to become de facto policy makers through proprietary systems that public institutions cannot meaningfully inspect or override.

9.7 Establish international red lines

Major AI powers should pursue agreements covering autonomous nuclear command, AI-enabled biological-weapons development, uncontrolled self-replication, attacks on frontier-AI safety infrastructure, and concealment of catastrophic incidents. Verification will be imperfect, but reciprocal monitoring, incident hotlines, compute and facility assurance, and shared attribution methods can raise the cost of violation without creating a single world AI authority.

9.8 Protect whistleblowers and independent scrutiny

Employees, contractors, researchers, journalists, and public officials need safe channels to report strategic-risk failures. Non-disclosure agreements should not suppress lawful reporting of catastrophic safety or illegal conduct. Oversight bodies require technical staff, secure access, subpoena or inspection authority, and protection from regulatory capture.

10. A Practical Governance Program

Horizon	Action	Purpose
Within 12 months	Create a strategic-capability map; identify single-vendor dependencies; mandate incident reporting; test continuity plans.	Discover present vulnerabilities before capability scales.
1–3 years	Adopt cross-domain merger review; require portability and redundancy; license catastrophic deployments; fund public-interest evaluators.	Prevent lock-in and unsafe coupling.
3–7 years	Build sovereign and allied countervailing AI; flexible biodefense manufacturing; resilient public communications; international red-line agreements.	Create durable defensive parity.
Ongoing	Publish annual concentration and dependency assessments; red-team successors and governance failure; revise thresholds based on	Keep safeguards institutionally neutral and technologically current.

	capability, not brand names.	
--	------------------------------	--

11. Avoiding Two Dangerous Extremes

Complete centralization is unsafe because one captured institution could control humanity's most powerful tools. Complete decentralization is also unsafe because unrestricted access can place catastrophic capabilities in the hands of irresponsible individuals, criminal groups, or covert state programs.

The safer equilibrium is distributed authority with controlled high-risk capability: multiple sovereign and private centers; genuine technical independence; reciprocal monitoring; constrained access to catastrophic functions; interoperability; and public capacity to intervene. The principle is simple: no one should control everything, and no one should be able to do everything merely because the technology exists.

12. Conclusion

The case for action does not require proving that Elon Musk, Peter Thiel, Palantir, Tesla, SpaceX, Anduril, or any associated network intends domination. Waiting for proof of intent would repeat a familiar governance mistake: regulating only after concentrated capability has been converted into harm.

The proper test is structural. Would society be comfortable if the same integrated capability were inherited tomorrow by an openly hostile successor? Could a democratic government continue essential operations if the provider withdrew cooperation? Could independent institutions see, understand, and stop misuse? If the answer is no, the concentration itself is a vulnerability.

Constitutional government did not emerge because every king was evil. It emerged because a political system dependent on a good king was recognized as fragile. The same insight now needs to be applied to artificial intelligence and the private infrastructures through which it will see, decide, communicate, and act in the world.

Humanity should never depend for its survival on the continuing benevolence, competence, or integrity of one AI, one corporation, one government—or one alliance. The goal is not fear of technological power. It is constitutionalism for technological power.

Sources and Further Reading

1. European Commission, “General-purpose AI obligations under the AI Act” (obligations applying from 2 August 2025; systemic-risk duties). [Source](#)
2. European Commission, “Guidelines on the scope of the obligations for providers of general-purpose AI models.” [Source](#)
3. Government of Argentina, “El Presidente Javier Milei recibió al empresario Peter Thiel,” 23 April 2026. [Source](#)
4. U.S. Defense Security Cooperation Agency, “Ukraine – Satellite Communications Services,” 29 August 2025. [Source](#)

5. UK AI Security Institute, “Frontier AI Trends Report 2025.” [Source](#)
6. OECD, “Steering AI’s Future: Strategies for Anticipatory Governance.” [Source](#)
7. U.S. Government Accountability Office, “Artificial Intelligence Acquisitions: Agencies Should Collect and Apply Lessons Learned to Improve Future Procurements,” GAO-26-107859. [Source](#)
8. Silicon Valley Defense Group, “NatSec100 Report: What Do the Companies Do?” (2026). [Source](#)

Analytical Note

This paper is a strategic-risk framework, not an accusation of unlawful conduct. Named organizations and individuals are included because their publicly documented activities illustrate the convergence problem. Assertions about future coordination, sovereign sanctuary, or hostile intent should be treated as hypotheses unless supported by specific, verifiable evidence.