

Public-Interest Surveillance and the Architecture of Democratic Risk

A refined research finding on how safety, health, efficiency, identity, and innovation narratives can enable durable surveillance infrastructures

AUTHOR	Samuel Abinsinguza · Founder, CentPol
SERIES	CentPol Papers No. 2
EDITION	Published 8 June 2026 · Version 1.0
TOPIC	Surveillance & democracy
ONLINE	www.centpol.com/research/white-papers/public-interest-surveillance

How public-interest justifications for surveillance reshape democratic risk – and the oversight architecture needed to govern it.

SUGGESTED CITATION

Abinsinguza, S. (2026, June 8). Public-Interest Surveillance and the Architecture of Democratic Risk (CentPol Papers No. 2). CentPol. <https://www.centpol.com/research/white-papers/public-interest-surveillance>

Executive finding

Modern surveillance capacity is rarely built in a single dramatic act. More often, it accumulates through systems introduced for plausible public purposes: reducing crime, improving emergency response, preventing fraud, distributing welfare, managing migration, protecting public health, easing travel, or modernising urban services. The research reviewed here is persuasive in identifying a recurring pattern: technologies adopted under benign or urgent rationales can become rights-threatening when they are made mandatory, linked across databases, retained indefinitely, governed opaquely, or made available to law-enforcement and intelligence bodies without meaningful public control.

The central finding should therefore be stated carefully. The problem is not that every biometric ID programme, public-health system, smart-city platform, data-sharing hub, or public-private technology contract is inherently illegitimate. Some forms of targeted surveillance can be lawful, necessary, and beneficial. The decisive question is whether the system is **legally bounded, necessary for a specific legitimate aim, proportionate to that aim, transparent to the public, technically constrained against reuse, independently audited, contestable by affected persons, and subject to deletion or sunset when the original justification expires**. Where these safeguards are absent, surveillance introduced as public service can harden into infrastructure for social control.

This refinement strengthens the original research in three ways. First, it replaces a polemical “play-book” framing with a more objective **risk-pattern analysis**. Second, it adds a clearer evidentiary standard: the report does not claim uniform intent by all governments or vendors; it identifies mechanisms through which capability, incentives, and weak oversight repeatedly produce surveillance expansion. Third, it anchors the analysis in recognised legal and policy standards, especially the human-rights principles of legality, necessity, proportionality, purpose limitation, non-discrimination, transparency, remedy, and independent oversight. The UN Human Rights Office has warned that networked digital technologies can become tools of “surveillance, control and oppression,” and has called for public-space surveillance to be restricted to measures that are strictly necessary and proportionate, with biometric recognition in public spaces subject to immediate restriction.¹

Evidence audit and refinement decisions

The original paper’s architecture is strong. Its seven categories capture real mechanisms: crisis justification, techno-solutionism, biometric identity anchoring, economic behavioural incentives, public-private data access, function creep, and data fusion. The main weakness is not the conceptual model, but the tone and evidentiary discipline. Phrases such as “true surveillance objective,” “manufacturing consent,” and “God’s-eye view” may be rhetorically effective, but they risk making the analysis appear predetermined. A research finding should sound less like an indictment and more like a careful explanation of how systems evolve under particular institutional conditions.

AREA REVIEWED	ASSESSMENT OF ORIGINAL DRAFT	REFINEMENT APPLIED
Core thesis	Substantively plausible, but too certain about intent across diverse political systems.	Recast as a conditional risk finding: surveillance becomes dangerous when public-interest systems lack binding limits and oversight.
Source quality	Strong civil-society sources are mixed with weaker sources such as Wikipedia, vendor blogs, and commercial listicles.	Replaced weak support with institutional and academic sources from OHCHR, HRW, Brennan Center, World Bank ID4D, Stanford SCCEI, Brookings, GDPR, the EU AI Act, and Cambridge University Press.
Objectivity	The original often treats benevolent justifications as inherently deceptive.	Added acknowledgement that public health, fraud prevention, welfare delivery, road safety, and crime prevention can be legitimate aims.
Causation	Some claims imply that adoption of technology automatically produces authoritarian control.	Clarified that risk depends on law, institutional context, interoperability, compulsion, retention, procurement incentives, and redress mechanisms.
Case studies	Case studies are relevant but some claims need qualification, especially around social credit, health-code repurposing, and Aadhaar integration with other systems.	Retained examples while narrowing claims to documented or carefully framed concerns.
Policy recommendations	Useful, but framed mainly as “resistance.”	Reframed as a governance test combining law, technical design, procurement controls, audit, remedies, and democratic authorisation.

Methodology and evidentiary standard

This finding is based on a qualitative synthesis of legal, human-rights, policy, and case-study literature. It uses comparative examples not to argue that all states follow a single hidden script, but to identify **recurring institutional mechanisms** that appear across different political settings. The standard of proof is therefore pattern-based rather than conspiratorial. A finding is treated as robust where multiple credible sources document the same mechanism, such as emergency powers becoming normalised, biometric systems creating exclusion risks, public-private data markets weakening warrant protections, or fusion centres expanding beyond their original counterterrorism mandate.

The analysis also distinguishes three separate questions that are often collapsed. The first is **purpose**: whether a surveillance measure pursues a legitimate public aim. The second is **capability**: what the system technically makes possible once deployed. The third is **governance**: whether law and institutions prevent the capability from being used for broader, discriminatory, or politically abusive ends. A programme may begin with a legitimate purpose and still create dangerous capability if governance is weak. This distinction is essential for objectivity.

The refined model: seven mechanisms of surveillance expansion

1. Crisis justification and emergency normalisation

Crises create the strongest political conditions for expanding surveillance. Terrorist attacks, pandemics, violent-crime spikes, migration emergencies, and national-security shocks can make expanded monitoring appear not merely useful but unavoidable. In such moments, the normal democratic demand for evidence, debate, judicial oversight, and sunset clauses is often treated as a luxury. The public may accept intrusions it would reject under ordinary conditions because the stated trade-off is immediate safety.

The post-9/11 United States remains a central example because it shows how emergency surveillance can become institutional architecture. The USA PATRIOT Act, the expansion of national-security authorities, and the development of fusion centres reflected a broader shift toward data-driven prevention. The Brennan Center notes that the U.S. fusion-centre network, initially promoted for counterterrorism information-sharing, grew to 80 centres and expanded into “all crimes and all hazards,” while suffering from weak oversight, flawed analysis, and monitoring of First Amendment-protected activity.² This does not mean that information-sharing after 9/11 was irrational. It means that emergency-born institutions require especially strict boundaries because their mandate tends to expand once personnel, databases, funding streams, and interagency habits are in place.

Public health presents a more complex case. Disease surveillance, contact tracing, and outbreak monitoring are legitimate and often necessary. The World Health Organization describes surveillance as an essential early-warning function for detecting and responding to public-health threats.³ Yet the COVID-19 period also showed how health-related digital systems can create durable identity, movement, and compliance infrastructures if they are not sunsetted, decentralised where appropriate, limited by purpose, and insulated from law-enforcement repurposing. The right lesson is not that public-health surveillance is illegitimate. It is that emergency health systems need public-law limits before they are deployed, not after they have become normal.

“Even where surveillance tools are initially rolled out for legitimate goals, they can easily be repurposed, often serving ends for which they were not originally intended.” – UN Human Rights Office, 2022.¹

2. Techno-utopian solutionism and the smart-city bargain

Smart-city and safe-city projects promise efficiency, sustainability, reduced crime, better transport, and improved delivery of public services. These aims are not trivial. Cities do face real problems, and data can help solve some of them. The risk emerges when urban management is redesigned around continuous sensing without democratic clarity about who collects the data, who analyses it, how long it is stored, whether it is shared with police or intelligence agencies, and whether residents can contest its use.

The “smart city” shifts surveillance from discrete investigation to ambient governance. Cameras, sensors, licence-plate readers, facial-recognition systems, transit cards, mobile-phone data, utility data, and predictive analytics can become a single operating layer for urban life. The citizen becomes legible as a moving data profile. In authoritarian settings, this can strengthen direct political control. In democratic settings, it can still produce discriminatory policing, opaque automated decision-making, and a loss of practical anonymity in public space.

Chinese surveillance exports illustrate both supply and demand. Brookings’ analysis of China’s global surveillance exports warns against simplistic explanations. Adoption is driven by “push” factors, including China’s geopolitical interests and the market power of Chinese technology firms, but also by “pull” factors in recipient states, including high crime rates and demand for public-security capacity.⁴ That nuance matters. Governments do not always buy these systems because they openly seek repression. They may buy them because they are under pressure to address violence, infrastructure failure, or administrative weakness. The rights risk is that once the technical system is installed, the same cameras and analytics that monitor traffic or crime can be used to monitor protests, opposition groups, journalists, or minority communities.

Human Rights Watch adds an important corrective: abusive surveillance should not be framed only as a China-versus-democracy problem. Similar surveillance technologies are used globally, and the difference between systems often lies in legal context, press freedom, civil society strength, judicial independence, and the degree of centralised state power.⁵ A rights-respecting smart city is therefore not defined by the brand of the vendor alone. It is defined by procurement transparency, independent impact assessment, strict purpose limitation, technical minimisation, public reporting, and enforceable remedies.

3. Biometric identity anchoring and the risk of exclusion

Biometric identity systems are often justified as tools for inclusion, anti-corruption, fraud reduction, and efficient service delivery. These are serious public aims. In countries without reliable civil registries, biometric deduplication can help ensure that one person does not hold multiple official identities and that public benefits reach intended recipients. The World Bank’s ID4D guidance recognises that biometrics can be accurate and efficient for deduplicating large populations, but it also states clearly that biometrics are not required or appropriate in all contexts and can create privacy, exclusion, cost, and operational risks.⁶

The core risk is that biometrics transform identity into an infrastructure of dependency. If a biometric ID becomes the gateway to food, schooling, banking, healthcare, mobile connectivity, taxation, travel, and welfare, then authentication failure is no longer a technical inconvenience. It can become social exclusion. This is especially acute for manual labourers with worn fingerprints, elderly people, persons with disabilities, people with visual impairments, children, people with albinism, and people in areas with poor connectivity or unreliable devices.⁶

India’s Aadhaar programme shows both the promise and the danger. It was promoted as a means to reduce welfare fraud and provide identity at enormous scale. Human Rights Watch and Amnesty International India warned, however, that making Aadhaar a prerequisite for essential services could

obstruct access to food, healthcare, education, and social security; they also documented concerns about biometric authentication failures, data breaches, broad deactivation powers, and weak grievance mechanisms.⁷ These concerns do not prove that all biometric ID is inherently abusive. They do show that biometric systems require legally guaranteed alternatives, human review, offline fallbacks, opt-out or limited-use options where feasible, independent security audits, and remedies that are fast enough to prevent denial of essential services.

A second biometric risk is interoperability. A biometric identifier becomes far more powerful when linked across welfare, banking, taxation, SIM registration, police databases, border systems, and facial-recognition networks. Even if each database is justified separately, the linked system can create a comprehensive identity graph. The safeguard is not merely cybersecurity. It is **structural separation**: different databases should not be linkable except under narrowly defined legal authority, independent authorisation, logged access, and strong penalties for misuse.

4. Economic scoring and behavioural governance

Economic surveillance becomes especially powerful when monitoring is connected to access: credit, licences, contracts, transport, schooling, insurance, employment, housing, or public benefits. In this form, surveillance is not only about watching. It is about shaping conduct through incentives and penalties. The most familiar example is China's social credit ecosystem, but the subject requires precision. Popular accounts sometimes imply a single universal score assigned to every citizen. The reality is more fragmented, varied across localities, and especially developed in corporate and regulatory contexts.

Stanford's analysis of China's Corporate Social Credit System describes a data-driven scoring system for rating the trustworthiness of business entities registered in China. It is linked to red lists and black-lists, regulatory rewards and sanctions, and collective enforcement across agencies. The research finds that politically connected firms in Zhejiang received higher scores partly by accumulating "soft merits" through party-state-sanctioned donations, volunteer activity, and government awards, raising concerns that the system can nudge firms toward policies preferred by the Chinese Communist Party.⁸ This is a stronger and more precise basis than broad claims about a single omnipotent social score.

The wider risk is that social-scoring logic can appear outside China in less centralised forms. Commercial data brokers, loyalty programmes, insurance telematics, fintech scoring, platform reputation systems, workplace productivity analytics, and predictive-risk tools can all convert behavioural data into opportunity allocation. A democratic society may not have a single state social-credit score, but it can still develop a patchwork of private and public scoring systems that produce similar consequences: people are sorted, priced, excluded, or targeted based on opaque profiles they cannot inspect or correct.

The EU AI Act's prohibited-practices framework is useful here because it identifies the danger at the level of mechanism rather than ideology. Article 5 prohibits AI social scoring that evaluates or classifies people over time based on social behaviour or personal traits where the result is detrimental treatment in unrelated contexts or treatment that is unjustified or disproportionate.⁹ This is the correct regulatory principle. The problem is not scoring as such. Credit scoring, risk assessment, and eligibility determina-

tion can be legitimate if narrow, transparent, accurate, contestable, and proportionate. The problem is generalised behavioural scoring that migrates across contexts and disciplines people for conduct unrelated to the decision being made.

5. Public-private entanglement and the data-broker loophole

The surveillance capacity of the modern state increasingly depends on private infrastructure.

Telecommunications firms carry communications; platforms host speech and social networks; cloud providers store records; brokers aggregate location and behavioural data; app developers collect sensor data; analytics companies convert raw data into predictions. This makes the boundary between state and corporate surveillance porous.

The legal problem is that constitutional and statutory safeguards often regulate direct government collection more clearly than government access to commercially available data. The Brennan Center argues that gaps in U.S. surveillance law permit the collection and use of Americans' communications and other protected information without adequate statutory authorisation or judicial oversight, including by purchasing data from brokers.¹⁰ This creates what might be called a procurement workaround: if the private sector has already collected the data, the state may claim it is merely buying information available in the marketplace rather than conducting a search.

This dynamic changes incentives. If government agencies become reliable customers for precise location trails, communications metadata, social-media monitoring, or behavioural profiles, the market has a reason to collect more granular data than consumer services require. Public-private entanglement therefore does not simply let the state access existing surveillance capitalism. It can help finance and normalise it.

A rights-based response must close the distinction between “collected by government” and “bought by government.” If the state could not lawfully compel production of data without a warrant, court order, or statutory process, it should not be able to evade that threshold by purchasing the same data.

Procurement law, privacy law, and constitutional doctrine must be aligned so that rights do not depend on whether surveillance is performed by a public official, a contractor, or a broker.

6. Function creep and incremental normalisation

Function creep is the gradual expansion of a system beyond its original purpose. It is among the most important mechanisms because it explains how surveillance can grow without a single visible moment of democratic choice. A system begins with a narrow aim: traffic enforcement, welfare deduplication, pandemic contact tracing, asylum processing, school safety, or crowd management. Over time, the data is retained longer, shared more widely, searched for new purposes, integrated with other databases, or upgraded with analytics that were not contemplated when the public first accepted the system.

The danger is not merely hypothetical. The UN Human Rights Office warns that surveillance tools initially introduced for legitimate goals can be repurposed.¹ In practice, function creep is encouraged by sunk costs. Once a city has installed cameras, sensors, command centres, and network infrastructure,

agencies can argue that new uses are efficient because the hardware already exists. Adding facial recognition, behavioural analytics, or protest monitoring may appear to be a software upgrade rather than a new political decision.

The antidote is legal and technical purpose limitation. GDPR Article 5 states that personal data must be collected for specified, explicit, and legitimate purposes; must be adequate, relevant, and limited to what is necessary; must not be kept in identifiable form longer than necessary; and must be processed with accountability.¹ Those principles are not bureaucratic formalities. They are anti-function-creep safeguards. Without them, every database becomes a future policing database, every camera becomes a future recognition device, and every emergency system becomes a standing administrative asset.

7. Data fusion and the inversion of suspicion

Data fusion is the capstone mechanism. It occurs when separate streams of information are pooled into a searchable intelligence environment: police records, travel data, financial transactions, immigration files, health records, CCTV feeds, licence-plate data, biometric identifiers, commercial data, and public tips. Fusion is attractive because it promises to “connect the dots.” The problem is that it can also dissolve the boundaries that protect ordinary life from continuous investigation.

The U.S. fusion-centre example shows the institutional pattern. Centres were established to improve information-sharing after 9/11, but their remit expanded from counterterrorism to all crimes and hazards. Brennan Center’s review found that they have produced flawed analysis, monitored First Amendment-protected activity, targeted minority communities and protest movements, and operated with inadequate federal oversight despite receiving federal funding, personnel, and database access.² The concern is not information-sharing per se. Investigators sometimes need to coordinate. The concern is generalised intelligence production about people not suspected of crime.

Data fusion can invert the presumption of innocence. Traditional investigation begins with a specific offence or credible suspicion and then seeks evidence under legal constraints. Fusion environments begin with broad data acquisition and search for suspicious patterns later. Suspicion becomes a product of the database rather than a predicate for accessing it. This is especially dangerous when inputs include unverified tips, biased policing records, commercial data of uncertain provenance, or algorithmic predictions that affected people cannot challenge.

Comparative governance context

The same surveillance technology has different consequences depending on political and legal context. A facial-recognition system deployed in a state with independent courts, adversarial media, strong data-protection authorities, public procurement transparency, and enforceable remedies is not equivalent to the same system deployed in a state where courts, media, police, and regulators are controlled by the ruling party. Yet democratic safeguards can erode, and authoritarian-style effects can appear in democratic systems when oversight becomes formal rather than real.

Human Rights Watch’s comparison of China and the United States is useful because it avoids moral simplification. It recognises the severity of Chinese state surveillance, especially in Xinjiang, while also noting serious rights abuses and surveillance capitalism in the United States. The key difference is not that one context has technology and the other does not. It is that legal and political constraints differ: courts, legislatures, press freedom, civil society, and fragmentation of power can mitigate abuse, but only if they remain effective.⁵

GOVERNANCE FACTOR	LOWER-RISK CONFIGURATION	HIGHER-RISK CONFIGURATION
Legal basis	Clear law defining purpose, scope, thresholds, retention, sharing, and remedies.	Executive discretion, secret rules, vague mandates, or emergency decrees.
Necessity and proportionality	Targeted use tied to a specific serious threat or administrative need.	Generalised population monitoring justified by broad safety or efficiency claims.
Purpose limitation	Data cannot be reused without new legal authorisation and public justification.	Data is freely repurposed across agencies once collected.
Technical architecture	Decentralised or minimised design, short retention, strong access logs, deletion by default.	Centralised databases, indefinite retention, broad searchability, weak audit trails.
Oversight	Independent regulator, courts, legislature, public reporting, and civil-society scrutiny.	Internal review only, classified procurement, weak parliamentary or judicial access.
Contestability	Individuals can know, challenge, correct, and obtain remedy.	Affected persons cannot see the data, score, source, or reason for adverse treatment.
Procurement	Human-rights impact assessment before acquisition; vendor transparency; ban on abusive use cases.	Vendor-driven adoption, opaque contracts, foreign financing, no rights assessment.

Strengthened case-study synthesis

India: democratic institutions and biometric dependency

India demonstrates that surveillance-risk architectures can emerge inside democratic systems, especially when identity infrastructure becomes central to public and private life. Aadhaar’s scale and ambition are unmatched: it aimed to provide identity, improve welfare delivery, and reduce fraud. Those aims are legitimate. The rights issue is that a single identifier connected to many services creates a dependency structure. When authentication fails, when data is leaked, when enrolment is difficult, or when deactivation occurs without adequate remedy, the harm is not abstract privacy loss but direct exclusion from food, banking, education, healthcare, and social protection.⁷

The Indian case also shows why data-protection law alone is not enough. A privacy statute can help, but surveillance risk depends on how identity systems interact with police databases, welfare systems, private authentication markets, mobile connectivity, and financial services. The strongest safeguard would

combine data-protection law with sector-specific limits on linkage, independent audit of authentication failures, mandatory alternatives for essential services, and strict prohibitions on using welfare identity infrastructure for generalised policing.

China: centralised party-state power, social management, and technological integration

China is the clearest case of surveillance technology operating in a political environment with limited independent constraints on state power. Surveillance in Xinjiang, public-space monitoring, censorship, platform control, and policing technologies form part of a wider governance model in which security, social stability, and party authority are tightly connected. HRW has documented how mobile apps, biometric collection, AI, and big data have been used to identify people deemed “untrustworthy,” particularly Uyghurs and other Turkic Muslims.⁵

At the same time, China should be analysed carefully rather than mythologically. The social credit system is not best understood as a single all-seeing score for every person. It is a set of evolving regulatory, financial, local, and corporate systems using data to enforce trustworthiness, compliance, and state priorities. Stanford’s work on the corporate system is especially valuable because it shows behavioural governance in an empirically grounded way: firms can be nudged toward party-state preferences through scoring categories, rewards, and sanctions.⁸

The United States: fragmented power, commercial data, and national-security exceptions

The United States is not a centralised surveillance state in the Chinese sense, but it has significant surveillance vulnerabilities. Post-9/11 authorities, intelligence programmes, fusion centres, predictive policing, protest monitoring, commercial data brokerage, and platform surveillance create a complex ecosystem in which rights may be undermined through fragmentation rather than centralisation. Brennan Center’s work on fusion centres and surveillance-law gaps shows how oversight can fail even where formal legal institutions exist.^{2 10}

The U.S. case matters because it prevents the analysis from becoming a story about authoritarian others. Democratic systems can generate surveillance through emergency law, local policing incentives, private markets, weak procurement controls, and judicial doctrines that lag behind technology. The rights danger is not only the authoritarian state that openly demands obedience. It is also the democratic state that quietly buys data, expands information-sharing, and allows private profiling to become public power.

Developing-country contexts: need, financing, and institutional asymmetry

Surveillance exports often succeed where governments face real governance pressures: violent crime, weak civil registries, limited administrative capacity, infrastructure gaps, and public demand for security. Brookings’ analysis rightly stresses that recipient-state demand is central; high-crime contexts may adopt Chinese surveillance platforms because officials believe these systems can solve urgent public problems.⁴ This matters because advocacy that simply warns governments not to buy foreign surveillance technology may fail if it does not address the underlying needs.

The more effective approach is to insist on rights-preserving alternatives: transparent procurement, public consultation, independent human-rights impact assessment, local data-residency and access rules, contractual bans on political monitoring, open auditing of accuracy and bias, and clear deletion schedules. Exporting states and vendors should not be allowed to shift rights risk onto weaker institutions. OHCHR has called for robust export-control regimes and human-rights impact assessments for surveillance technologies, including assessment of both technical capability and conditions in the recipient country.¹

The governance test: when public-interest surveillance becomes dangerous

A surveillance measure should be treated as high-risk when several conditions converge. It collects sensitive or biometric data at scale; participation is mandatory or practically unavoidable; the data is retained beyond the immediate purpose; agencies can search or share it without independent authorisation; private vendors can access or monetise it; affected persons cannot know or challenge adverse use; and there is no sunset clause, deletion rule, or independent audit. Under those conditions, the stated purpose may remain benign while the system's actual governance profile becomes dangerous.

The EU AI Act offers a useful model for the highest-risk AI-enabled surveillance practices. Article 5 prohibits certain forms of manipulative AI, exploitative systems, social scoring, untargeted facial-image scraping, and some biometric categorisation. It also restricts real-time remote biometric identification in publicly accessible spaces for law-enforcement purposes to narrowly defined exceptions, requiring strict necessity, proportionality, fundamental-rights assessment, registration, and prior judicial or independent administrative authorisation.⁹ These rules reflect a broader principle: some systems are so structurally prone to abuse that they require prohibition or near-prohibition rather than ordinary compliance paperwork.

For systems that are not banned, governance must be built into design. A rights-respecting system should collect the least data necessary, avoid centralised identifiers where possible, separate administrative databases from policing access, use short retention periods, publish procurement contracts and impact assessments, record every access event, allow independent audit, provide meaningful notice, and guarantee human review and remedy. The burden should be on the deploying authority to prove necessity and proportionality, not on citizens to prove future harm.

Policy recommendations

For policymakers and legislatures

Legislatures should enact comprehensive data-protection and surveillance-governance laws that apply to both direct government collection and government purchase of commercially available data. The law should require clear statutory authority for surveillance technologies, define permissible purposes narrowly, impose retention limits, prohibit incompatible secondary use, and require independent authori-

sation for access to sensitive data. GDPR Article 5's principles of lawfulness, transparency, purpose limitation, data minimisation, storage limitation, security, and accountability provide a concise baseline, but surveillance law must go further for public authorities and police access.¹¹

Public bodies should not acquire high-risk surveillance technologies through ordinary procurement channels. Before acquisition, agencies should publish a surveillance impact assessment, a human-rights impact assessment, technical documentation, accuracy and bias testing, cost-benefit analysis, retention rules, sharing rules, and alternatives considered. Local democratic approval should be required for municipal systems such as facial recognition, licence-plate readers, gunshot detection, predictive policing, or smart-city sensor networks. Secret pilots should be prohibited except in genuinely narrow national-security contexts subject to independent classified oversight.

Certain systems should face moratoriums or bans. Real-time facial recognition in public spaces, untar- geted scraping of facial images to create recognition databases, generalized social scoring, emotion in- ference in schools or workplaces, and predictive criminality assessments based solely on profiling are incompatible with democratic rights unless restricted to exceptionally narrow and independently au- thorised circumstances. The EU AI Act's prohibited-practices framework is a useful legislative reference point.⁹

For courts and oversight bodies

Courts and regulators should reject formalistic distinctions between data compelled by government and data bought from brokers. If a dataset reveals intimate patterns of movement, association, religion, health, sexuality, politics, or communication, the method of acquisition should not determine the level of protection. Oversight bodies should have technical capacity, subpoena power, access to classified or confidential procurement materials, and authority to suspend systems that violate legal limits.

Independent audits should examine not only cybersecurity, but also disparate impact, false positives, false negatives, mission creep, access logs, data-sharing chains, vendor compliance, and actual opera- tional value. A system that is invasive but ineffective fails proportionality. Surveillance cannot be justi- fied by aspiration; it must demonstrate necessity in practice.

For civil society, journalists, and researchers

Civil society should focus on the chain of deployment: budget proposals, vendor demonstrations, pro- curement contracts, pilot programmes, data-sharing agreements, retention schedules, and memoranda of understanding. The most consequential decisions often happen before a technology becomes visible. Public-records requests, litigation, technical audits, community education, and coalition building re- main essential because surveillance systems are frequently justified in technical language that excludes ordinary democratic participation.

Advocacy should avoid claiming that every safety or identity technology is inherently oppressive. That claim is easy for governments to dismiss. A stronger approach is to ask precise questions: What problem does this solve? What evidence shows it works? What less intrusive alternative was considered? What data is collected? Who can access it? How long is it retained? Can it be used for law enforcement? Can affected people challenge the result? What happens when the emergency ends?

For technology companies and investors

Companies should conduct human-rights due diligence before selling surveillance-relevant systems, especially to police, intelligence, border, migration, or public-security agencies. Due diligence should evaluate not only the direct customer but also the political context, rule-of-law environment, intended use, foreseeable misuse, data access arrangements, and downstream integration. Vendors should refuse deployments that enable indiscriminate public monitoring, political repression, or discriminatory targeting. Transparency reports should disclose government contracts, data requests, model limitations, accuracy testing, and known rights risks wherever legally possible.

Investors and insurers can also shape the market. They should treat high-risk surveillance as a governance and liability issue, not merely a growth sector. Systems that cannot be audited, explained, or constrained should be considered legally and reputationally unstable assets.

Conclusion

The original research is directionally sound: public-good narratives can help normalise surveillance infrastructures that later become difficult to challenge. Its strongest insight is that surveillance expansion is cumulative. A crisis creates urgency. A smart-city project installs sensors. A biometric ID anchors identity. A public-private contract supplies data. A pilot expands through function creep. A fusion centre combines the streams. A scoring or eligibility system turns observation into consequence.

The refined finding is more cautious but also more durable. It does not require proving that every actor intends repression from the start. It is enough to show that **capabilities, incentives, and weak safeguards interact predictably**. Once built, surveillance systems seek new uses. Once databases exist, agencies seek access. Once vendors find a market, they expand collection. Once emergency measures become routine, the burden shifts from the state justifying surveillance to the public trying to dismantle it.

The democratic answer is not technological rejection. It is democratic control. Surveillance power must be made specific, temporary, minimised, reviewable, contestable, and accountable. A society can use technology for safety and public service without accepting permanent population monitoring as the price of modern life. But that outcome is not automatic. It has to be designed, legislated, audited, and defended before the infrastructure becomes too ordinary to see.

References

1. United Nations Office of the High Commissioner for Human Rights, "Spyware and surveillance: Threats to privacy and human rights growing, UN report warns," 16 September 2022, <https://www.ohchr.org/en/press-releases/2022/09/spyware-and-surveillance-threats-privacy-and-human-rights-growing-un-report>.
2. Mike German, Rachel Levinson-Waldman, and Kaylana Mueller-Hsia, "Ending Fusion Center Abuses," Brennan Center for Justice, 15 December 2022, <https://www.brennancenter.org/our-work/policy-solutions/ending-fusion-center-abuses>.
3. World Health Organization, "Surveillance in emergencies," <https://www.who.int/emergencies/surveillance>.
4. Sheena Chestnut Greitens, "Dealing with demand for China's global surveillance exports," Brookings Institution, April 2020, <https://www.brookings.edu/articles/dealing-with-demand-for-chinas-global-surveillance-exports/>.

5. Maya Wang, Frederike Kaltheuner, and Amanda Klasing, "The future of technology: Lessons from China—and the US," Human Rights Watch, 9 May 2023, <https://www.hrw.org/news/2023/05/09/future-technology-lessons-china-and-us>.
6. World Bank Identification for Development, "Biometric data," ID4D Practitioner's Guide, <https://id4d.worldbank.org/guide/biometric-data>.
7. Human Rights Watch and Amnesty International India, "India: Identification Project Threatens Rights," 13 January 2018, <https://www.hrw.org/news/2018/01/13/india-identification-project-threatens-rights>.
8. Stanford Center on China's Economy and Institutions, "China's Corporate Social Credit System and Its Implications," summarising Lauren Yu-Hsin Lin and Curtis J. Milhaupt, "China's Corporate Social Credit System: The Dawn of Surveillance State Capitalism?," <https://sccei.fsi.stanford.edu/china-briefs/chinas-corporate-social-credit-system-and-its-implications>.
9. European Union Artificial Intelligence Act, Article 5, "Prohibited AI Practices," <https://artificialintelligenceact.eu/article/5/>.
10. Elizabeth Goitein, "How to Fix U.S. Surveillance Law," Brennan Center for Justice, 18 July 2023, <https://www.brennancenter.org/our-work/analysis-opinion/how-fix-us-surveillance-law>.
11. General Data Protection Regulation, Article 5, "Principles relating to processing of personal data," <https://gdpr-info.eu/art-5-gdpr/>.
12. Ana Beduschi, "Rethinking digital identity for post-COVID-19 societies: Data privacy and human rights considerations," *Data & Policy* 3, Cambridge University Press, 14 July 2021, <https://www.cambridge.org/core/journals/data-and-policy/article/rethinking-digital-identity-for-postcovid19-societies-data-privacy-and-human-rights-considerations/OB9A65B889C341CF535E804256C2816A>.