

Surveillance is built in increments.

Public-interest surveillance and the architecture of democratic risk.

Samuel Abinsinguza · Founder, CentPol

Surveillance capacity rarely arrives in one dramatic act.

It accumulates through systems introduced for plausible public purposes — reducing crime, improving emergency response, preventing fraud, distributing welfare, protecting public health, modernising urban services.

Technologies adopted under benign or urgent rationales can become rights-threatening when they are made mandatory, linked across databases, retained indefinitely, governed opaquely, or opened to law-enforcement and intelligence bodies without meaningful public control.

The finding is stated carefully. Not every biometric ID programme, public-health system or smart-city platform is illegitimate, and some targeted surveillance is lawful, necessary and beneficial. *The decisive question is governance.*

A system is defensible only when it is:

- ✓ **legally bounded** and necessary for a specific legitimate aim;
- ✓ **proportionate** to that aim and transparent to the public;
- ✓ **technically constrained** against reuse;
- ✓ **independently audited** and contestable by the people it affects; and
- ✓ subject to **deletion or sunset** when its original justification expires.

A pattern, not a conspiracy

The paper refines earlier research that framed surveillance as a "playbook". It replaces polemic with **risk-pattern analysis**: no claim of uniform intent by governments or vendors, but recurring mechanisms through which capability, incentives and weak oversight produce surveillance expansion across very different political systems. A finding counts as robust where multiple credible sources document the same mechanism.

It also separates three questions that are often collapsed: **purpose** — does the measure pursue a legitimate aim; **capability** — what the system makes technically possible once deployed; and **governance** — whether law and institutions stop that capability being used for broader, discriminatory or politically abusive ends. A programme can begin with a legitimate purpose and still create dangerous capability when governance is weak.

Seven mechanisms of surveillance expansion

1 • Crisis

Attacks, pandemics and crime spikes make expanded monitoring look unavoidable, and evidence, debate and sunset clauses get treated as luxuries. U.S. fusion centres, built for counterterrorism, grew to 80 and spread to "all crimes and all hazards".

2 • Smart city

Continuous sensing — cameras, plate readers, transit and phone data — turns surveillance from discrete investigation into ambient governance. The risk depends less on the vendor's brand than on procurement transparency, purpose limits and remedies.

3 • Biometric ID

Deduplication can serve inclusion and anti-fraud aims, but when one identifier gates food, banking, schooling and welfare, an authentication failure becomes exclusion — and linked databases become an identity graph.

4 • Scoring

Monitoring tied to access — credit, licences, contracts, benefits — shapes conduct through rewards and penalties. Outside China it appears as a patchwork of private and public scores people cannot inspect or correct.

5 • Data brokers

Safeguards often bind direct collection more tightly than purchase. Buying commercially available data becomes a procurement workaround — and a market signal to collect more.

6 • Function creep

Systems built for one aim are retained longer, shared wider and upgraded with analytics nobody approved. Sunk costs make each new use look like a software update rather than a political decision.

7 • Data fusion

Pooling police, travel, financial, health, biometric and commercial data inverts suspicion: it becomes a product of the database rather than the predicate for searching it.

The mechanisms compound. 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.

Context decides the consequence

The same facial-recognition system is not equivalent in a state with independent courts, adversarial media and enforceable remedies and in one where courts, media, police and regulators answer to a ruling party. Yet democratic safeguards erode, and authoritarian-style effects appear in democracies, when oversight becomes formal rather than real.

FOUR CASES

India shows how a single identifier connected to many services creates dependency inside a democracy: when Aadhaar authentication fails, the harm is exclusion from food, banking and education. **China** is the clearest case of surveillance with limited independent constraint — and its social credit system is better understood as evolving regulatory, local and corporate systems than as one all-seeing score. **The United States** shows surveillance through fragmentation: fusion centres, protest monitoring and data brokerage. In **developing countries**, demand is real — violent crime, weak registries, thin administrative capacity — so advocacy that simply says "don't buy" fails unless it offers rights-preserving alternatives.

The governance test

Treat a measure as high-risk when several conditions converge: sensitive or biometric data at scale; participation that is mandatory or practically unavoidable; retention beyond the immediate purpose; search or sharing without independent authorisation; vendor access or monetisation; no way for affected people to know or challenge adverse use; and no sunset, deletion rule or independent audit.

Some practices are so prone to abuse that the EU AI Act prohibits them outright — social scoring, untargeted facial-image scraping — and restricts real-time remote biometric identification in public spaces to narrow, pre-authorised exceptions. For everything else, governance must be built into design: least data necessary, separation of administrative databases from policing access, short retention, published contracts and impact assessments, logged access, independent audit, and human review with remedy.

The burden belongs on the deploying authority to prove necessity and proportionality — not on citizens to prove future harm.

Recommendations

Legislatures: one standard for data the state collects and data it buys; clear statutory authority, narrow purposes, retention limits and independent authorisation; no high-risk surveillance through ordinary procurement; local democratic approval for systems such as facial recognition and plate readers; moratoriums on real-time public facial recognition, generalised social scoring and emotion inference in schools and workplaces.

Courts and oversight bodies: reject the line between compelled and purchased data; audit disparate impact, error rates, mission creep and actual operational value — an invasive but ineffective system fails proportionality.

Civil society and journalists: follow the deployment chain — budgets, vendor demonstrations, pilots, data-sharing agreements — and ask precise questions rather than claiming every safety technology is oppressive. **Companies and investors:** run human-rights due diligence on the customer's context, refuse indiscriminate or political monitoring, and treat unauditable systems as unstable assets.

The democratic answer is not technological rejection. It is democratic control: surveillance power made specific, temporary, minimised, reviewable, contestable and accountable.

About CentPol

CentPol is a **global community of people — from wildly different professional backgrounds — curious enough to understand, shape, and build the future of AI** rather than watch it happen to them.

We produce **policy intelligence for the next technology era** through community, stakeholder dialogue, and research. This paper is part of CentPol's research on how societies can use technology for safety and public service without accepting permanent population monitoring.

This brief summarizes *Public-Interest Surveillance and the Architecture of Democratic Risk* (CentPol Papers No. 2) by Samuel Abinsinguza, published 8 June 2026.

Read the full white paper.

The evidence audit, the governance-risk table, case studies from India, China, the United States and developing countries, and recommendations by audience — with 12 references.

centpol.com/research/white-papers/public-interest-surveillance