Scientific Research Publishing

Beijing Law Review, 2026, 17(4), 1144-1166
https://www.scirp.org/journal/blr
ISSN Online: 2159-4635
ISSN Print: 2159-4627

Who Owns the Smart City? Smart Cities, Invisible Walls

Cesar Beck

School of Law, ATITUS University, Passo Fundo, Brazil
Email: ocesarbeck@gmail.com

How to cite this paper: Beck, C. (2026). Who Owns the Smart City? Smart Cities, Invisible Walls. *Beijing Law Review, 17,* 1144-1166.
https://doi.org/10.4236/blr.2026.174056

Received: August 18, 2026
Accepted: November 7, 2026
Published: November 10, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/

Open Access

Abstract

The smart city was built around a persuasive promise: technology would make urban government work better. As cities have become more data-intensive, however, that promise has acquired a less familiar counterpart. The same infrastructures that help governments understand traffic, services, or public space also make the people moving through those environments increasingly knowable. This article examines the institutional conditions under which AI-enabled urban surveillance can serve legitimate public purposes without allowing informational power to escape meaningful public control. Methodologically, it combines doctrinal legal analysis, interdisciplinary smart-city scholarship, contemporary regulatory materials, and a case study of Flock Safety's automated license-plate recognition network in the United States. The analysis develops **asymmetric visibility** as a condition in which technological systems acquire increasing knowledge about individuals while those individuals possess comparatively limited knowledge of, access to, or control over the infrastructures observing them. The Flock case shows that legitimate public-security capabilities can become institutionally dependent on privately operated informational infrastructure, while the American data-broker ecosystem demonstrates why information mobility beyond its original purpose also requires legal attention. The central finding is institutional rather than technological: the same capability can be defensible in one governance arrangement and difficult to justify in another. The article responds by asking law to follow urban information through its institutional life: why the city needs it, how much visibility that purpose can justify, where the information may travel, who answers for its use, and what remains open to challenge. These questions form the Public-Interest Smart Governance Framework proposed here.

Keywords

Smart Cities, Artificial Intelligence, Urban Surveillance, Asymmetric Visibility, Privacy, Public Accountability

1. Introduction

A car passes through an American street. No officer stops it; no checkpoint interrupts the trip. To the driver, nothing has happened. Yet the car may already have left a record behind: plate, location, time, vehicle characteristics. An automated license-plate reader can preserve what would otherwise have been an unremarkable passage through an intersection. Multiply that event across a network that Flock Safety says now serves more than 5000 law-enforcement agencies and 6000 communities, processing more than 20 billion plate reads each month, and the object begins to change. These are the company's numbers and should be read as such. Still, their scale is difficult to dismiss. If the plate was already visible in public, what changed? Persistence, searchability, and networked access can turn an observation that once would have vanished into part of an informational infrastructure (Flock Safety, n.d.).

That infrastructure can serve legitimate ends. There are obvious reasons why police departments want such a tool. ALPR systems can assist investigations, help recover stolen vehicles, and locate missing persons; even Oxnard, while suspending its Flock deployment in 2026, continued to describe the technology as highly valuable. That tension is worth keeping in view. The smart city did not enter public debate as a surveillance project. It promised something considerably more ordinary and attractive: services that worked better, mobility with less friction, greater connectivity, and new ways of participating in urban life (Hollands, 2008; Townsend, 2013). That promise remains evident in contemporary institutional policy. ITU's 2025 work on smart sustainable cities places AI and IoT among tools for resilient, inclusive, and sustainable urban services, while the OECD treats AI as a means of improving transport, planning, and other public functions while warning about privacy, bias, and accountability (ITU, 2025c; OECD, 2025a).

The complication is that connection is no longer the whole story. A connected city exchanges information. A datafied city also remembers, combines, and searches it. Once ordinary movement becomes machine-readable, the question shifts from whether technology reaches the citizen to what the resulting architecture can know about the citizen. Earlier work used the image of *technological walls* to describe unequal access to smart-city benefits (Beck, Boff, & Cenci, 2022). The problem now runs in both directions. Digital exclusion persists, but a person can be poorly served by digital public infrastructure and still be highly legible to systems that observe vehicles, devices, or movement patterns.

That imbalance is what this article calls asymmetric visibility: a condition in which technological systems acquire increasing knowledge about individuals, while those individuals possess comparatively limited knowledge of, access to, or control over the infrastructures that observe them. This is partly a privacy problem. Not quite only that. It is also a distributive and institutional problem: who receives the benefits of urban intelligence, who bears its scrutiny, who can understand the systems involved, and who can challenge them when something goes wrong?

The institutional problem sharpens when public authorities rely on privately designed and operated infrastructure. That reliance is not inherently suspect. Private firms may supply expertise or scale that public agencies cannot efficiently reproduce. Nor does public ownership guarantee good governance. The harder question is whether a public institution retains meaningful legal and practical control when observation, storage, search, permissions, and interoperability depend upon a vendor's platform. Then ownership and control begin to pull apart.

Recent American experience makes the problem concrete. In August 2025, an official Illinois audit found that U.S. Customs and Border Protection had obtained access through Flock's network to Illinois license-plate camera data in violation of state restrictions; the Secretary of State ordered the access shut off, and Flock subsequently paused federal pilot programs nationwide (Illinois Secretary of State, 2025). In February 2026, Oxnard suspended 19 Flock ALPR cameras after an audit found that two nationwide federal queries had included the city's data despite its *California-only* settings. The city stated that Flock accepted responsibility for the platform error (City of Oxnard, 2026). Neither episode proves that privately operated surveillance infrastructure is inherently abusive. They show something narrower: a legal boundary can exist on paper and still fail in the technical architecture through which public power is exercised.

The research question follows from that tension: under what institutional conditions can AI-enabled urban surveillance serve legitimate public purposes without turning the smart city into privately governed infrastructure for networked monitoring? For this article, AI-enabled urban surveillance refers to interconnected technological systems that collect, retain, retrieve, and analyze information about activities in urban environments, potentially incorporating artificial intelligence to identify patterns or generate inferences. These functions should not be conflated. Automated license-plate recognition (ALPR) converts vehicle observations into machine-readable records; data storage and networked search enable their subsequent retrieval and circulation; and AI-supported inference may extend their analytical significance. Not every ALPR deployment necessarily involves advanced artificial intelligence, nor does the legal significance of persistent surveillance depend exclusively on algorithmic inference.

My hypothesis is institutional rather than technological. The same capacity may be legitimate when tied to a defined public purpose, lawful authority, necessity, proportionality, and effective accountability, but it becomes much harder to justify when operational control shifts to opaque or weakly governed private infrastructure. The technology has not necessarily changed. The institutional conditions around it have.

This article examines how AI-mediated urban environments alter the relationship among governance, privacy, inequality, and public responsibility, focusing on privately operated surveillance infrastructure in the United States. Methodologically, the article combines doctrinal legal analysis, interdisciplinary smart-city scholarship, and qualitative case analysis. Its legal and institutional materials in-

clude U.S. federal and state regulatory provisions, governmental audits, administrative enforcement actions, and policy documents, with particular attention to California and Illinois. The Flock Safety incidents were selected because they illustrate discrepancies between legally established access restrictions and their technical implementation, while the Federal Trade Commission's data-broker enforcement actions provide a distinct perspective on the secondary circulation of location information. These cases are examined comparatively to identify recurring institutional vulnerabilities involving data access, information mobility, and public accountability, rather than to establish statistically representative findings. Time-sensitive factual claims are, where possible, grounded in primary governmental and regulatory materials from 2025-2026; company sources are treated as evidence of what the company reports or how it describes its own systems, not as independent evidence of effectiveness.

The argument moves through three stages. Chapter 1 follows the smart city from connectivity to datafication and develops asymmetric visibility from that transition. Chapter 2 asks what happens institutionally when useful State capacity comes to depend on private informational infrastructure, using Flock and the American data-broker market for different parts of that inquiry. Chapter 3 turns from diagnosis to law and develops the Public-Interest Smart Governance Framework. By then, the title's question should look somewhat different. Ownership matters, but control over how the city sees, remembers, and uses information matters more.

2. From Smart Cities to Invisible Walls

2.1. The Promise of the Smart City

More technology was never, by itself, a convincing theory of the smart city. A street does not become smarter because another sensor has been attached to a lamppost. The claim was always more demanding: digital infrastructure should solve recognizable urban problems and make ordinary life in the city work somewhat better. Mobility, public services, environmental management, and participation were the point. The technology was supposed to earn its place by serving them.

That distinction has accompanied the idea from the beginning. Townsend (2013) described a city in which physical space and digital systems were becoming increasingly intertwined. Hollands (2008), writing from a more critical position, warned against assuming that technological sophistication makes a city socially progressive. The tension between those perspectives remains useful. A city can acquire considerable computational capacity and still govern badly. The number of connected devices is therefore a poor measure of smartness unless those devices improve something that matters to the people living among them.

Contemporary policy has not abandoned that people-centered ambition. The ITU's 2025 work situates connectivity, infrastructure, AI, and the Internet of Things within a broader project of resilient and inclusive urban development ra-

ther than treating technological deployment as an achievement in itself (ITU, 2025c). The OECD arrives at much the same point from another direction: a city is not made *smart* by the mere presence of digital systems, but by what those systems contribute to well-being, public services, sustainability, and inclusion (OECD, 2025a).

There is an obvious practical reason for the attraction. The pressure behind this technological turn is not imaginary. The ITU reports that 57% of the world's population was already urban in 2022 and projects that 68% will be by 2050 (ITU, 2025c). At that scale, administrative ignorance has costs of its own. Traffic, energy, infrastructure, and public safety generate patterns that rarely respect the boundaries between municipal departments. Better information cannot solve those problems by itself, but governing them while deliberately knowing less is hardly a serious alternative.

So the promise is still worth taking seriously. The difficulty lies elsewhere. The technologies through which cities now pursue efficiency and responsiveness are becoming capable of producing far more information about urban life than the earlier language of connectivity alone suggests. The smart city still wants to work better for its inhabitants. Increasingly, it does so by knowing more about what happens around them.

2.2. From Connectivity to Urban Datafication

Connectivity remains part of the architecture, but it is no longer the destination. Linking an agency to a database or a traffic system to a network allows information to move. Contemporary urban systems increasingly add another capacity: they convert events in the physical city into information that can be stored, retrieved, combined, and analyzed. The connected city has become a data-producing city. The vocabulary used by ITU-T Study Group 20 for 2025-2028 is revealing in itself: IoT sits beside digital twins, edge computing, AI, machine learning, big data, and advanced sensing (ITU, 2026). These are not interchangeable technologies, and treating them as such would obscure more than it explains. Their relevance here lies in what happens when they meet. An event can be sensed, transmitted, preserved, connected to other records, and later interpreted. What once disappeared into the ordinary noise of city life can become available for another question at another time.

The OECD describes the same development from the perspective of urban use. The OECD's examples show what this looks like in practice. Urban AI is increasingly used not only to organize what a city already knows about transport, planning, energy, or public services, but also to ask what may happen next: where demand might rise, how a system might respond, or what a simulated change would produce (OECD, 2025a). Digital twins make the distinction particularly visible because their utility depends on maintaining computational representations of physical environments through continuing flows of information.

The change is easier to grasp at street level. Someone waiting to cross an avenue

may watch hundreds of cars go by. By dinner, most are forgotten; a week later, almost certainly all of them are. That ordinary act of forgetting once placed a natural limit on observation. A database has no comparable need to forget. An observation that seemed irrelevant when it occurred can remain searchable after its significance changes. Add observations from other places or other moments, and something else becomes possible: not merely remembering an event, but relating it to others.

That is why aggregation matters. One record can answer where something was at a particular moment. A sufficiently rich sequence may reveal a pattern. The same intuition extends beyond vehicles. The same logic travels across the city. A transit tap, an energy reading, or a camera image may say little on its own. Its informational character changes when it can be retrieved months later, compared with other records, or connected to another system. Storage keeps the observation alive. Interoperability lets it travel. Analysis can make it say more than it said when first collected.

The implications extend beyond the technical capabilities of these systems. Brayne (2017), examining the incorporation of large-scale data infrastructures into policing, shows how information once scattered across institutional records can become part of an integrated surveillance apparatus. Joh (2016) approaches the problem from a legal perspective, examining how automated surveillance and data-driven policing reshape investigative discretion. Their analyses help explain why the significance of urban surveillance cannot be reduced to initial information collection. What matters is also what institutions can subsequently retrieve, connect, and act upon. A record that appears inconsequential today may acquire an entirely different significance when placed alongside information collected tomorrow.

None of this makes datafication inherently objectionable. In fact, many of the benefits claimed for contemporary smart cities depend on precisely these capabilities. The OECD's 2025 work points to applications in transport optimization, infrastructure, energy, and urban planning, while also emphasizing that privacy, bias, transparency, and accountability become more important as AI moves deeper into city systems. The legal shift is subtler. Once cities increasingly understand and manage themselves through persistent data, the informational representation of urban activity begins to function as infrastructure in its own right. And that creates a question connectivity alone did not have to answer: what happens when becoming part of the smart city also means becoming legible to it?

2.3. From Technological Exclusion to Asymmetric Visibility

Datafication does not make the older digital divide disappear. The numbers resist any easy declaration that the access problem has been solved. In 2025, the ITU estimated that roughly 6 billion people were online and another 2.2 billion were still offline. Geography sharpened the divide: Internet use reached 85% among urban residents but 58% among rural residents, falling to only 14% among rural residents in low-income countries (ITU, 2025b). Connectivity may be ordinary

for some urban inhabitants; it remains scarce for many others. A city does not solve technological inequality merely by becoming technologically sophisticated. But exclusion no longer describes the whole problem. Consider the apparently paradoxical position of someone with poor access to digital public services who nevertheless moves every day through streets equipped with cameras, sensors, and connected infrastructure. That person may remain under-connected to the benefits of the digital city while being perfectly reachable by systems that record aspects of urban activity. Technological disadvantage and technological visibility are not opposites. They can exist at the same time.

The technological wall has therefore acquired another side. One wall limits the citizen's access to the system. An invisible one may allow the system to reach the citizen without offering comparable knowledge of what happens afterward. What was collected? How long will it remain? Can another institution search it? Can separate observations be connected? The difficulty lies not simply in being observed, but in inhabiting an informational environment whose architecture is far easier for institutions to see than for the individual moving through it. I use asymmetric visibility to describe this imbalance. The asymmetry is not simply that the system sees more. It is that the two sides occupy radically different positions regarding what has been seen. The institution may know what was collected, how records connect, who can retrieve them, and how long they remain useful. The person moving through the city may know almost none of this. Nor do collection and opacity necessarily rise together: an information-rich system can be comparatively transparent, while a narrower one may leave practically no route for understanding or challenging its actions.

Is this simply a privacy problem? Not quite. Privacy remains indispensable, but it does not quite capture the distributional problem. The benefits of a datafied city and the burden of being visible to it need not fall on the same people. Some inhabitants may understand the systems around them, possess resources to challenge mistakes, or encounter relatively little technological scrutiny. Others may not. The point is not that every neighborhood or person should be observed to precisely the same degree. Different places present different traffic patterns, risks, and public needs. Unequal visibility is not necessarily illegitimate. Unjustified asymmetric visibility may be.

The invisible wall, then, is not defined only by who remains outside the technological city. It can also appear among those already inside it, through unequal positions of knowledge, exposure, and control. That leaves a different question for the next stage of the argument. If greater urban visibility can serve legitimate public purposes, who controls the infrastructure that makes such visibility possible—and what changes when that infrastructure is no longer entirely the State's own?

3. When Public Safety Becomes Private Infrastructure

3.1. Urban AI as State Capacity

Government cannot administer a complex city without knowing something about

it. Governments have always sought to render populations and territories administratively legible. Registers, maps, traffic records, and infrastructure monitoring are hardly inventions of artificial intelligence. What AI changes is the practical ceiling: far more information can be compared, processed, and acted upon before human administrative capacity becomes the bottleneck. It does not create the underlying governmental need to know. Some of that additional capacity is plainly useful. The OECD's study of 200 government AI use cases found applications in public services, decision support, forecasting, and anomaly detection (OECD, 2025b). Urban government makes the attraction particularly obvious. Traffic changes while the paperwork is still moving. Floodwater is indifferent to departmental jurisdiction. A missing person can cross a municipal boundary in minutes. In cases like these, computational speed is not simply administrative convenience; it can become part of the State's ability to act while action still matters.

Emergency management provides a useful example. Current ITU work describes the use of digital twins, IoT, and AI to integrate geospatial information, sensor feeds, and other operational data for public safety and disaster resilience. Its Singapore case describes an integrated digital environment designed to support emergency simulations and coordinated response. Here, greater technological visibility is not an unfortunate side effect of the system. Information is what makes the public function possible (ITU, 2025a).

Public safety raises the same point more sharply. If technology can help locate a missing person, recover a stolen vehicle, or reconstruct events relevant to a serious crime, then its informational capacity has public value. A technologically capable State is not, for that reason alone, a surveillance State. But usefulness settles less than it first appears to. The OECD's own assessment complicates the optimistic picture. Bad data do not become better because they move faster; opacity does not become accountability because a model is efficient; and automation can propagate an error as readily as it propagates a useful decision (OECD, 2025b). A system can therefore perform its technical function successfully while still raising legal questions about purpose, reach, retention, or responsibility.

The choice is not between an informationally blind government and one entitled to know everything technology makes knowable. The harder question is how much visibility a particular public function actually requires. That question will become central later. For now, the baseline is simpler: technologically enhanced State capacity can be legitimate, and private firms can legitimately help provide it. The institutional difficulty begins one step later. What happens when the State does not merely purchase a device, but comes to depend on a privately operated system to collect, retain, search, or connect the information through which the public function is performed? At that point, the question is no longer only what the government can see. It is who controls the infrastructure that allows it to see.

3.2. Flock Safety and the American ALPR Network

Flock Safety is useful as a hard case because the public value of its technology is

easy to understand. Its automated license-plate readers capture plate information and vehicle characteristics, turning those observations into records that authorized users can later search. Police agencies use ALPR systems in investigations involving stolen vehicles, violent crimes, and missing persons. Flock's own 2025 Impact Census, based on responses from roughly 700 law enforcement agencies and company data, estimates that its technology supported more than one million investigations or incidents and helped locate more than 10,000 missing persons during the year (Flock Safety, 2026c). These are corporate estimates, not independent measurements of effectiveness. They nevertheless help explain why public agencies want the technology.

At some point, however, scale changes the object we are describing. Flock says its network reaches more than 5000 law-enforcement agencies across 49 states and more than 6000 communities, processing over 20 billion plate reads per month (Flock Safety, n.d.). Those are corporate figures, not independent measurements, but they tell us what the company itself claims to have built. One camera records a passing vehicle. A network makes that observation useful somewhere else. The legal significance lies increasingly in that second fact.

The platform also contains safeguards that should not be ignored. Flock states that customers own their data and control access and sharing; its current published policy sets a 30-day default retention period, subject to local law or contractual variations, after which data are deleted. The company says that its LPR systems do not use facial recognition and that it does not sell customer data. Searches are tied to individual users, require an investigative purpose, and create audit records that supervisors can review (Flock Safety, 2026b). Under the company's current policy, sharing is opt-in rather than automatic.

Those controls matter because network effects operate in both directions. Cross-jurisdictional access may be helpful when a stolen vehicle or a missing person crosses municipal boundaries. The very feature that increases investigative usefulness, however, also increases the legal significance of permissions. An isolated local database has one informational boundary. A networked platform may expose that boundary to decisions made elsewhere. Illinois showed what happens when the legal and technical boundaries diverge. In August 2025, an audit by the Illinois Secretary of State found that U.S. Customs and Border Protection had accessed Illinois license-plate camera data through Flock's system in violation of state restrictions governing ALPR information sharing under the Illinois Vehicle Code, 625 ILCS 5/2-130(b)-(c) (Illinois, 2026). The audit began with a sample of twelve local law enforcement agencies and concluded that adequate safeguards had not prevented the access. It also uncovered a Flock pilot with CBP that, according to the Secretary of State, Flock leadership did not know about. Illinois ordered the access to be shut off, and Flock subsequently paused its CBP and other federal pilots nationwide (Illinois Secretary of State, 2025). The problem was not that Illinois lacked a legal boundary. The platform failed to reproduce it reliably.

Oxnard exposed a related, but slightly different, failure. The California police

department had deployed 19 fixed Flock ALPR cameras and had disabled National Lookup from the outset. California (2026) Civil Code § 1798.90.55(b) restricts public agencies from selling, sharing, or transferring ALPR information except to another public agency, and only as otherwise permitted by law. This statutory restriction should be distinguished from Oxnard's California-only platform configuration, which represented an additional operational limitation on network access. After learning of a vendor-related problem affecting California agencies, Oxnard audited more than five million queries received by its system during 2025. It found that two nationwide searches by the U.S. Department of Veterans Affairs Police Department in May 2025 had included Oxnard's data despite the city's California-only settings. The searches concerned criminal investigations rather than immigration enforcement—a qualification worth preserving because the problem was not the motive behind those particular queries. It was access itself. Oxnard reported that Flock accepted full responsibility for the platform error, and the city suspended the cameras until it could be satisfied that the restrictions it had chosen would actually hold (City of Oxnard, 2026). That qualification matters. Oxnard did not conclude that ALPR was useless or inherently illegitimate. Its police chief called the technology highly valuable even as it suspended the deployment. The dispute concerned the architecture through which a legitimate investigative capacity was governed. Organizational intention said one thing. The system produced another.

Flock has since modified that architecture. In March 2026, the company acknowledged that some California camera networks had been inadvertently accessible to out-of-state agencies during 2025 and said that limitations in earlier technical logging made the precise cause impossible to establish in some cases. It introduced additional California-specific guardrails and had already developed search filters, more structured search reasons, and audit tools (Flock Safety, 2026a).

The architecture is still changing. As of the company's currently published policy materials, thirty days remain the default retention period. On August 13, 2026, however, Flock announced a further set of reforms, including the mandatory use of its Audit Assistance system, stronger links between searches and official cases, more granular controls over inter-agency sharing, and a planned reduction of the standard retention window from 30 days to 7 (Rankin & Lauer, 2026). Contemporary reporting indicates that several of those requirements are to become mandatory by the end of 2026 or January 2027, rather than having already taken effect (Rankin & Lauer, 2026).

The timing is revealing. These changes do not prove that every safeguard will work, just as the earlier failures do not prove that networked ALPR cannot be governed responsibly. They demonstrate something more useful for this article: permissions, retention, auditing, and sharing are not peripheral features of the technology. They are part of its governance architecture.

That is why formal data ownership does not settle the institutional question. A

police department may own the records generated by its cameras while relying on a private platform to manage permissions, organize searches, maintain audit trails, and connect information across agencies. The public function remains public. Part of the informational machinery through which it becomes possible does not. This is the arrangement the next section describes as privatized urban intelligence.

3.3. Privatized Urban Intelligence and the Accountability Gap

Flock matters beyond license-plate recognition because it exposes a more general problem. Governments buy technology all the time. A private company supplying software, cameras, or cloud services does not become problematic merely because it is private, nor does procurement somehow cease to be public administration. The harder distinction is functional: when does the vendor cease to be merely a supplier and become part of the infrastructure through which the public function itself is performed?

This is no longer a problem that can be left to procurement theory. In April 2026, the U.S. Government Accountability Office reported that federal agencies had more than doubled their use of AI between 2023 and 2024 and were still acquiring additional capabilities through fiscal year 2025. Crucially, the government was not always dictating the technological path: the GAO found both agency-directed acquisitions and vendor-driven ones. Agencies bought AI both as products and as continuing services, while officials identified difficulties involving technical expertise, performance monitoring, data rights, and testing requirements (GAO, 2026). The OECD has identified a related risk in public procurement: proprietary technology and data formats can create vendor or data lock-in, leaving public authorities increasingly dependent on a supplier they may find costly or technically difficult to replace (OECD, 2025b).

Private participation is not the threshold I am interested in. Dependence is. A city can buy a camera, a database, or a software service and remain perfectly capable of governing the public function behind it. Something changes when the function begins to depend on an architecture that the authority cannot readily replace, fully inspect, or operate without the vendor. I call that condition privatized urban intelligence: public informational capacity whose practical operation depends substantially on privately designed or maintained infrastructure. There is no single diagnostic fact. Proprietary architecture, control over updates, custody of data, interoperability, and switching costs matter because, taken together, they tell us how difficult it would be for the State to recover practical control.

Dependence also rearranges responsibility. The police department authorizes a search; the municipality signs a contract; the vendor builds the permissions; someone else may host the data. Lawyers can separate those relationships neatly enough on paper. The person affected by the system encounters their combined effect. That is the **accountability gap** at issue here—not an absence of legal responsibility, which would be too strong a claim, but the growing difficulty of finding and

exercising it when the technology behaves as one system while responsibility is divided among several institutions.

Contracts consequently do more than allocate commercial risk. They can shape retention, access, audit rights, interoperability, technical responsibilities, and even the government's ability to migrate away from a system. The citizen affected by those choices is usually not a party to the agreement. Yet those contractual decisions may determine important features of the technological environment through which public authority reaches that citizen.

Code adds another layer. Lessig's familiar insight that technological architecture can regulate conduct becomes unusually concrete here (Lessig, 1999). Law may prohibit a form of access. A contract may allocate responsibility for preventing it. But the software still has to make the prohibition operational. Illinois and Oxnard exposed that problem precisely: the legal or organizational boundary existed; the technical architecture did not reliably reproduce it.

Law, contract, and code, therefore, need not say the same thing. When they align, private technological capacity may reinforce public administration without displacing public control. When they diverge, formal authority can remain public while part of the operational logic through which it acts becomes private. That is the outsourcing paradox in its narrowest form. The function remains public; some of the machinery that makes it possible does not. Outsourcing may change where technological control sits. It should not, by itself, determine where public responsibility ends. The remaining problem concerns the information itself. Even if responsibility for the original public function can be located, what happens when data generated in one context acquires value somewhere else?

3.4. Data Brokers and the Secondary Market for Visibility

The institutional problem does not end with the system that first collects information. Data can acquire value elsewhere. That proposition requires a careful distinction. Nothing examined here shows that Flock operates as a data broker or sells its customers' ALPR records; the company states the opposite. The American data-broker market enters the argument for a different reason. It provides documented evidence that location information generated in one setting can later circulate among different actors and serve purposes that were not apparent when the underlying activity occurred. Flock illustrates privately mediated public infrastructure. Data brokers illustrate information mobility. The two should not be conflated.

Federal enforcement makes the latter problem unusually concrete. Mobilewalla shows how far an apparently ordinary location trace can travel. In January 2025, the Federal Trade Commission finalized an order against the company after alleging that it had tracked and sold sensitive location information without taking reasonable steps to verify consent. The resulting data revealed visits to places where the location itself says something intimate about a person—including healthcare facilities and places of worship (FTC, 2025a). The FTC separately finalized an or-

der against Gravy Analytics and its subsidiary, Venntel, after alleging that they had unfairly sold sensitive location information and collected or used location data without verifiable consent for commercial and governmental purposes (FTC, 2025b).

The important point is not that every stream of location data follows this route. It plainly does not. What these proceedings establish is narrower: information generated through ordinary activity can retain value after leaving its original context, and that value can support uses quite different from the one a person may have associated with generating the data in the first place.

California's DROP system is revealing for a reason that has little to do with the portal's mechanics. Since August 1, 2026, registered brokers must periodically retrieve centralized deletion requests through the platform, with the CPPA providing a 45-day window for the first batch (CPPA, 2026). The State has effectively created a single door through which a person can address companies they may never have known held their data. That institutional design tells us something about the market it regulates: information has become easier to circulate than its holders are to identify.

Secondary market for visibility is a useful shorthand for this phenomenon, but it need not become another major concept. The legal point is simpler. Collection is not the end of the data lifecycle, and remaining in the same database is not the only way information can escape its original purpose. A new user, a new query, or a new institutional connection may be enough. That closes the institutional problem developed in this chapter. Flock shows how a legitimate public function can depend upon private informational infrastructure. The data-broker cases show why boundaries around information continue to matter after collection. The question that follows is therefore not whether data should ever move. Legitimate cooperation may require that they do. It is when movement beyond the original purpose requires a new legal justification.

4. Governing the Invisible City

4.1. From the Right to the City to the Right to the Smart City

The right to the city offers a useful place to begin because its concern has never been technology as such. At its core lies a more basic question: whether inhabitants can meaningfully access the opportunities, services, and common resources that urban life produces. UN-Habitat has framed the concept in similar distributive terms, emphasizing equal access to urban resources and opportunities, particularly for groups otherwise excluded from the advantages of city life (United Nations Human Settlements Programme, 2020).

A datafied city complicates that familiar question. Access still matters, but it no longer captures the full relationship between the inhabitant and the infrastructure. A person may use public transport, cross a monitored street, enter a public facility, or simply move through the city while generating information that can later be stored, searched, or connected. Urban technology therefore operates in both di-

rections: it gives the citizen access to services, while sometimes giving the system access to the citizen.

None of this requires inventing a new fundamental right for each new layer of technology. That temptation should probably be resisted. Law already possesses much of the vocabulary it needs; the harder task is making that vocabulary work in an urban environment increasingly governed through data. Privacy and data protection deal with the processing of information. Equality asks who enjoys the benefits and who carries the burdens. Proportionality limits the reach of public power, while administrative accountability keeps the question of responsibility alive when that power operates through several institutions at once. The right to the city gives these commitments an urban setting. The problem is not the absence of legal concepts. It is whether familiar concepts can still do their job when part of the city has become informational.

The current UN-Habitat draft framework points in much the same direction. Its people-centered approach refuses to treat digital infrastructure as an achievement in itself: technology matters insofar as it improves urban life while remaining compatible with human rights, inclusion, and the needs of groups that technological modernization can otherwise leave behind (United Nations Human Settlements Programme, 2025). There is a consequence here that is easy to miss. For years, much of the smart-city debate has asked whether citizens can access the digital city. A datafied city forces the question to run the other way as well. Under what conditions should the digital city gain access to them? The right to the smart city cannot concern only the citizen's access to technology; it must also concern technology's access to the citizen.

The Sustainable Development Goals offer a useful triangulation rather than another framework to be superimposed on the argument. SDG 11 keeps the city itself in view: inclusion, safety, resilience, and sustainability. SDG 10 prevents those ambitions from being detached from inequality. SDG 16 adds the institutional question that becomes unavoidable once technological capacity is distributed across public and private actors: who remains accountable? Taken together, the three goals do not solve the problem identified here. They do make its structure harder to ignore. A city may distribute the benefits of information in one direction while distributing scrutiny, control, and responsibility quite differently.

The right to the city gets us this far, but not much further. Streets, housing, transport, and public services remain obvious objects of urban law. Data, permissions, and technical architectures are less visible, although they increasingly decide how those same spaces and services are administered. They belong to the city too. The remaining task is consequently practical: to identify the legal boundaries that allow urban intelligence to remain useful without letting the capacity to know expand simply because the technology permits it.

4.2. The Public-Interest Smart Governance Framework

By this point, neither extreme is particularly convincing. A city that refuses useful

informational tools may become less capable of protecting or serving its inhabitants. A city that adopts them whenever they promise efficiency has simply moved the difficult questions downstream. Illinois and Oxnard already showed where some of those questions reappear: in permissions, sharing, and responsibility. The legal problem is therefore not whether visibility should exist. It is what must constrain it when it does.

Existing AI-governance instruments reject the idea that risk can be addressed in a single moment. NIST's AI Risk Management Framework distributes the problem across governance, mapping, measurement, and management, with purpose and context remaining relevant throughout the system's lifecycle (NIST, 2023). The OECD reaches a similar practical conclusion regarding government use of AI: institutional capacity is valuable, but it must coexist with oversight, transparency, and responsibility (OECD, 2025b). This is not a case for freezing technological development. It is a recognition that the conditions under which a system operates are part of what makes its use defensible.

The Council of Europe's Framework Convention makes it harder to treat the point as merely managerial. Privacy, data protection, oversight, accountability, and procedural safeguards are legal concerns, not just desirable properties of a well-designed system. The Convention is particularly relevant here because it also reaches AI activities carried out by private actors on behalf of public authorities (Council of Europe, 2024). The public-private boundary, in other words, does not make the rights problem disappear.

Urban infrastructure makes those general commitments unusually concrete. A city has to decide why it needs to observe in the first place. It then has to decide how much observation the task actually requires, what may happen to the information afterward, who answers for a system partly operated by private actors, and where an affected person can turn when something goes wrong. None of these questions was invented to populate a five-part framework. Each emerged from the problems already encountered in the cases above.

I organize the answer around five legal anchors: Public Purpose, Proportional Visibility, Data Containment, Public Accountability, and Contestability. The names matter less than the sequence of problems behind them. Why is the city looking? How much does it need to see? What happens to the information afterward? Who still answers when private infrastructure sits in the middle? And, finally, what can be done when the system's output has consequences for someone? The Public-Interest Smart Governance Framework is an attempt to keep those questions connected.

The framework also requires institutional consequences when these principles cannot be demonstrated. Under Public Purpose, the responsible public authority must establish a sufficiently specific and lawful objective before authorizing deployment; otherwise, procurement or implementation should not proceed. Proportional Visibility requires the deploying agency to justify the scope of collection, retention, searchability, and access, with unjustified capabilities subject to re-

striction or removal. Data Containment places responsibility on the public authority and its technology provider to ensure that information-sharing permissions reflect applicable legal and contractual boundaries; where these safeguards fail, access must be restricted and continued operation reassessed. Public Accountability requires the contracting authority to preserve effective oversight through enforceable obligations concerning audits, incident reporting, corrective measures, and vendor cooperation. Finally, Contestability requires the responsible public institution to maintain legally appropriate avenues for reviewing consequential uses of surveillance information, including administrative or judicial remedies where applicable. These responsibilities may be distributed operationally, but their fragmentation cannot eliminate the public authority's ultimate obligation to justify and govern the exercise of public power.

There is nothing especially utopian in that proposal. Sometimes a city really does need to know more in order to govern well. The difficulty begins when the ability to know grows faster than the legal reasons that justify it. The framework is intended to keep those two things from drifting apart.

4.3. Public Purpose and Proportional Visibility

Public safety is the kind of justification that becomes less useful the more expansive it is. Almost any system of urban observation can be attached to it at a sufficiently high level of abstraction. Public Purpose demands something more specific. What is the authority trying to do, and why does doing it require this information? Locating a stolen vehicle gives law something concrete to assess. *Improving security* gives it much less. Specificity does not settle legality, but without it, there is no stable point from which necessity, retention, access, or later use can be judged.

NIST approaches risk in much the same way: the intended purpose, deployment context, applicable law, expected benefits, and potential impacts are part of the assessment of the system itself (NIST, 2023). That matters here because efficiency is impossible to evaluate legally in the abstract. A system can be extraordinarily efficient at doing something the public authority was never sufficiently justified in doing.

Purpose gets us only so far. A legitimate objective can still be pursued through an architecture that sees too much. The harder question is therefore not simply whether the State may observe, but how visible a person must become for the State to perform the task at hand. Proportional Visibility treats that degree of legibility as part of the proportionality inquiry. The justification must extend beyond the initial act of collection. It must also cover the informational reach built around it.

Proportional Visibility is not simply data minimization under another name. Counting records tells us how much a system collected; it may tell us surprisingly little about what the system can know. One plate read *places a vehicle somewhere*. A searchable sequence begins to describe movement. Connecting that sequence to other records changes the object again. The proportionality inquiry should follow

that change. Duration, retrievability, connection, and inference matter because each can make the same underlying data reveal more than they did when first collected.

The distinction is evident in the Flock case without requiring condemnation of the technology itself. Locating a stolen vehicle or investigating a serious offense can provide a legitimate reason to search license-plate observations. It does not follow that the same purpose automatically justifies every retention period, every cross-jurisdictional search, or every subsequent use. Illinois and Oxnard are useful precisely because public authorities had attempted to draw boundaries around access. The controversy arose when the technical architecture did not reliably preserve them.

Nor does Proportional Visibility require an informationally blind government. That would defeat many of the legitimate functions discussed in Chapter 2. Different public problems may also justify different levels of observation. The principle asks for something narrower: justification for each meaningful increase in informational reach. Why this information? Why for this long? Why searchable by these actors? Why connected to these systems?

Public Purpose and Proportional Visibility therefore perform different work. The first establishes why the city needs to see. The second asks how much the city needs to see to do what it has lawfully set out to do. A technologically capable State can answer both questions. Indeed, its legitimacy increasingly depends on being able to do so.

4.4. Data Containment, Public Accountability, and Contestability

A legitimate purpose and proportionate visibility still leave one question unanswered: where can the information go? Data Containment begins there. Information collected for one public purpose should remain institutionally and functionally bounded unless another use has its own lawful justification. This does not mean forcing every dataset into isolation. A stolen vehicle, missing person, or emergency may cross municipal borders, and useful public systems must sometimes do the same. The point is narrower: interoperability is a technical capacity, not legal permission.

Chapter 2 showed why the distinction matters. Illinois and Oxnard involved information becoming accessible beyond boundaries public authorities had attempted to impose. The data-broker cases demonstrated a different form of mobility, in which location information acquired value beyond its original context. Data Containment addresses both problems at the level of principle. Collection is only the beginning of the data lifecycle. New access, sharing, retention, or downstream use may require a new justification.

The next question is responsibility. Private infrastructure can divide operational control among agencies, vendors, and technical providers, but fragmentation cannot serve as a route out of public law. Procurement, contractual safeguards, oversight, and remedies remain governmental concerns when the under-

lying technology is used to exercise a public function. The principle can be stated plainly: the State may outsource technology; it cannot outsource public accountability.

This difficulty is hardly unique to urban surveillance. Citron (2008) demonstrated how automated administrative systems can weaken procedural safeguards when institutional oversight fails to keep pace with technological implementation. Brauneis and Goodman (2018), examining algorithmic transparency in smart cities, identified similar obstacles arising from government procurement practices and limited access to information about privately developed systems. Both contributions expose a recurring institutional problem: public authorities may retain formal responsibility for decisions while losing practical control over the technologies that shape them. A contract may allocate obligations between an agency and its vendor, but it cannot, by itself, resolve the question of accountability to the public.

That rule becomes important precisely when responsibility is difficult to locate. An agency may attribute a failure to the vendor; the vendor may point to customer configuration; the contract may again allocate obligations differently. To the affected person, however, there was one consequential exercise of public power. The Council of Europe's Framework Convention requires accountability and responsibility for AI-system activities and provides for documentation, complaint mechanisms, and procedural safeguards, while the NIST AI RMF calls for defined roles, monitoring, and mechanisms through which relevant actors can provide feedback or seek redress (Council of Europe, 2024; NIST, 2023).

That leads to Contestability. If technologically mediated information contributes to an investigation, identification, classification, or other consequential exercise of public authority, some meaningful avenue for review must remain available. Depending on context, that may involve correction, audit, administrative review, complaint procedures, or judicial remedies. Contestability does not require public disclosure of every investigative technique. Legitimate secrecy can coexist with review. What matters is that secrecy does not become practical immunity from it.

Oxnard offers a useful test of how these five anchors operate together. Under Public Purpose, the city's use of ALPR technology for legitimate criminal investigations may be justified, but that purpose does not automatically authorize unrestricted access by other agencies. Proportional Visibility requires the scope of searchable information and the range of authorized users to remain connected to the investigative objective. Data Containment becomes decisive when searches extend beyond the California-only restrictions established by the city: the relevant question is not merely whether those searches served lawful investigative purposes, but whether the information was accessible under the applicable legal and institutional conditions. Public Accountability then requires the deploying agency to investigate the failure, demand corrective measures from its vendor, and verify the effectiveness of access controls before restoring operations. Finally, Contesta-

bility requires appropriate mechanisms through which consequential uses of improperly accessed information may be examined or challenged, subject to applicable procedural safeguards. Oxnard's decision to suspend its cameras illustrates an important institutional response, but suspension alone does not establish that every governance concern has been resolved. The framework therefore shifts the inquiry from whether surveillance technology is useful to whether its actual operation remains legally justified, technically constrained, and institutionally answerable.

The lesson extends well beyond Oxnard. A city may have sound reasons to adopt surveillance technologies, but those reasons cannot excuse failures in how the systems are designed, operated, or supervised. Nor can public authorities treat contractual arrangements or technical complexity as convenient explanations for decisions they remain responsible for making. The challenge is not to make cities less intelligent, but to ensure that their growing capacity to observe does not outpace their capacity to justify. After all, a smart city should be measured not only by what its technology allows it to know, but also by what its institutions are prepared to answer for.

5. Conclusion—Who Owns the Smart City?

The smart city remains an attractive project for good reasons. Cities need information to manage mobility, infrastructure, emergencies, public safety, and increasingly complex services. Digital systems can make public institutions faster, more responsive and, in some circumstances, more capable of solving problems that conventional administration handles poorly. Nothing in this article requires abandoning that promise.

What has changed is the informational architecture through which it is pursued. Connectivity once captured much of what made a city *smart*. It no longer does. Contemporary urban systems can observe, retain, connect, and search information generated by ordinary activity. That capacity changes the inhabitant's position. Digital exclusion remains real, but access and exposure no longer move together: a person may receive few of the benefits of technological modernization while becoming highly legible to the infrastructures governing urban life. Asymmetric visibility names that imbalance.

The evidence examined here supports the central hypothesis, although with an important qualification. The decisive variable is neither whether a technology is public or private, nor whether it is observable. Its legitimacy depends on the institutional conditions under which that capacity becomes operational. Flock makes this unusually clear because ALPR can serve legitimate purposes. Finding a missing person or locating a stolen vehicle is not a trivial benefit. Yet Illinois and Oxnard showed that once observation becomes networked infrastructure, legal boundaries must also survive in permissions, search functions, and technical architecture. The problem was not that public authorities had no rules. In important instances, the systems through which those rules had to operate failed to repro-

duce them reliably.

The data-broker cases reveal a different part of the same institutional problem. They do not establish a commercial relationship with Flock, and none has been assumed here. They demonstrate something narrower: location information can retain value beyond the context in which it was generated. Collection is therefore only one moment in the governance of urban information. What happens afterward matters too.

The answer, then, cannot be to make the State informationally blind. Nor is there anything inherently illegitimate about private firms supplying the infrastructure through which public institutions become more capable. The difficulty begins when capability starts supplying its own justification: because information can be collected, it is collected; because systems can be connected, they are connected; because a record can be retained, somebody eventually finds another reason to use it. Law has to interrupt that logic. Public authority needs reasons for seeing, limits on what follows from seeing, and responsibility that does not disappear when the machinery sits elsewhere. The framework proposed here is an attempt to make those limits part of technological capacity rather than an afterthought.

This also answers the question in the title. Who owns the smart city? Formal ownership is almost the wrong question. A municipality may own data, a police department may operate a system, a vendor may control software, another provider may host information, and contractual permissions may determine what different institutions can actually do. What matters more is who exercises effective control over the informational architecture through which the city sees, remembers, searches, and acts upon information about its inhabitants—and whether that control remains subordinate to law and public responsibility.

The right to the smart city must therefore concern both directions of the technological relationship. Citizens should be able to access the opportunities created by digital urban development. They should also remain protected when the digital city reaches back. The contemporary invisible wall may not exclude a person from a street, a service, or a platform. It may instead be embedded in databases, permissions, contracts, and systems that determine who can see, who can know, and who can challenge. A city does not cease to build walls simply because they can no longer be seen.

Declaration of Generative AI and AI-Assisted Technologies in Manuscript Preparation

During manuscript preparation, the authors used Claude (Anthropics) solely for limited language editing of selected sentences in order to revise in order to improve clarity. No AI tool was used to generate research content, analyze data, interpret results, or prepare references, the ideas, argument, and intellectual contribution of the paper are entirely original. All AI-assisted suggestions were reviewed, revised where necessary, and approved by the authors, who take full re-

sponsibility for the final manuscript.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

Beck, C. A. M. R., Boff, M. M., & Cenci, D. R. (2022). Cidades inteligentes: Desigualdades, gentrificação e os desafios da implementação dos ODS. *Revista de Direito Econômico e Socioambiental, 13,* 565-593. https://doi.org/10.7213/revdireconsoc.v13i3.29005

Brauneis, R., & Goodman, E. P. (2018). Algorithmic Transparency for the Smart City. *Yale Journal of Law & Technology, 20,* 103-176.

Brayne, S. (2017). Big Data Surveillance: The Case of Policing. *American Sociological Review, 82,* 977-1008. https://doi.org/10.1177/0003122417725865

California (2026). *California Civil Code § 1798.90.55(b): Automated License Plate Recognition Systems*. California Legislative Information. https://leginfo.legislature.ca.gov/faces/codes_displaySection.xhtml?sectionNum=1798.90.55.&lawCode=CIV

California Privacy Protection Agency (CPPA) (2026). *Information for Data Brokers: Delete Request and Opt-Out Platform (DROP)*. California Privacy Protection Agency. https://cppa.ca.gov/data_brokers/

Citron, D. K. (2008). Technological Due Process. *Washington University Law Review, 85,* 1249-1313.

City of Oxnard (2026). *News Release: Oxnard Police Department Suspends Use of Flock Safety Automated License Plate Readers*. City of Oxnard. https://www.oxnard.gov/pd-news/news-release-oxnard-police-department-suspends-use-of-flock-safety-automated-license-plate-readers-02-27-26

Council of Europe (2024). *Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225)*. Council of Europe. https://www.coe.int/web/conventions/full-list?module=treaty-detail&treatynum=225

Federal Trade Commission (FTC) (2025a). *FTC Finalizes Order Banning Mobilewalla from Selling Sensitive Location Data*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-finalizes-order-banning-mobilewalla-selling-sensitive-location-data

Federal Trade Commission (FTC) (2025b). F*TC Finalizes Order Prohibiting Gravy Analytics, Venntel from Selling Sensitive Location Data*. Federal Trade Commission. https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-finalizes-order-prohibiting-gravy-analytics-venntel-selling-sensitive-location-data

Flock Safety (2026a). *Flock Implements Enhanced Guardrails across California to Ensure Lawful and Responsible Use of LPRs*. Flock Safety. https://www.flocksafety.com/blog/flock-implements-enhanced-guardrails-across-california-to-ensure-lawful-and-responsible-use-of-lprs

Flock Safety (2026b). *How Does Flock Handle License Plate Data Deletion*? Flock Safety. https://www.flocksafety.com/blog/how-does-flock-handle-license-plate-data-deletion

Flock Safety (2026c). *How Effective Is Flock? 2025 Impact Census Results from 700 Law Enforcement Agencies*. Flock Safety. https://www.flocksafety.com/blog/how-effective-is-flock

Flock Safety (n.d.). *Corporate Security*. Flock Safety. https://www.flocksafety.com/industries/corporate-security

Hollands, R. G. (2008). Will the Real Smart City Please Stand Up? Intelligent, Progressive or Entrepreneurial? *City, 12,* 303-320. https://doi.org/10.1080/13604810802479126

Illinois (2026). *Illinois Vehicle Code, 625 ILCS 5/2-130: User of Automated License Plate Readers; Prohibitions*. Illinois General Assembly.

Illinois Secretary of State (2025). *Giannoulias' Audit Finds License Plate Reader Company in Violation of State Law*. Illinois Secretary of State. https://www.ilsos.gov/news/2025/august-25-2025-giannoulias-audit-finds-license-plate-reader-company-in-violation-of-state-law.html

International Telecommunication Union (ITU) (2025a). *Citiverse Use Case Taxonomy: Public Safety, Health and Disaster Resilience*. International Telecommunication Union. https://www.itu.int/hub/publication/t-gen-meta-2025-6/

International Telecommunication Union (ITU) (2025b). *Measuring Digital Development: Facts and Figures 2025*. International Telecommunication Union. https://www.itu.int/hub/publication/d-ind-ict_mdd-2025-3/

International Telecommunication Union (ITU) (2025c). *Question 1/2: Smart Sustainable Cities and Communities*. International Telecommunication Union. https://www.itu.int/hub/publication/d-stg-sg02-01-2-2025/

International Telecommunication Union (ITU) (2026). *ITU-T Study Group 20: Internet of Things, Digital Twins and Smart Sustainable Cities and Communities (Study Period 2025-2028)*. International Telecommunication Union. https://www.itu.int/en/ITU-T/studygroups/2025-2028/20/Pages/default.aspx

Joh, E. E. (2016). The New Surveillance Discretion: Automated Suspicion, Big Data, and Policing. *Harvard Law & Policy Review, 10,* 15-42.

Lessig, L. (1999). *Code and Other Laws of Cyberspace*. Basic Books.

National Institute of Standards and Technology (NIST) (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. U.S. Department of Commerce.

Organisation for Economic Co-Operation and Development (OECD) (2025a). *Artificial Intelligence for Advancing Smart Cities: Issues Note for the 5th OECD Roundtable on Smart Cities and Inclusive Growth*. OECD. https://www.oecd.org/en/events/2025/10/5th-oecd-roundtable-on-smart-cities-and-inclusive-growth.html

Organisation for Economic Co-Operation and Development (OECD) (2025b). *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. OECD Publishing.

Rankin, S., & Lauer, C. (2026). *Flock Announces Changes amid Backlash over Its License Plate Reader Network*. AP News. https://apnews.com/article/flock-license-plate-cameras-surveillance-deflock-2a93bc075e2f7ffcca9e04a35d75a3fe

Townsend, A. M. (2013). *Smart Cities: Big Data, Civic Hackers, and the Quest for a New Utopia*. W.W. Norton & Company.

United Nations Human Settlements Programme (2020). *World Cities Report 2020: The Value of Sustainable Urbanization*. UN-Habitat. https://unhabitat.org/wcr/2020/

United Nations Human Settlements Programme (2025). *Draft International Guidelines on People-Centred Smart Cities: Smart Cities for People—Empowering Lives, Protecting*

the Planet, Advancing Prosperity (HSP/EB.2025/INF/10). UN-Habitat. https://unhabitat.org/meetings/first-session-of-2025-of-the-executive-board

United States Government Accountability Office (GAO) (2026). *Artificial Intelligence Acquisitions: Agencies Should Collect and Apply Lessons Learned to Improve Future Procurements (GAO-26-107859)*. U.S. Government Accountability Office. https://www.gao.gov/products/gao-26-107859