Journal of Electronic Research and Application, 2026, Volume 10, Issue 5
http://ojs.bbwpublisher.com/index.php/JERA
Online ISSN: 2208-3510
Print ISSN: 2208-3502

# Encoding Necessity for Standing Government Access to Platform-Held Data: A Three-Axis Model

**Linglan Xia***

J.S.D, Law School, Tianjin Normal University, Tianjin, China, 300382

***Corresponding author:** Linglan Xia, xllxia777@163.com



**Abstract:** Government access to platform-held data is increasingly implemented not through isolated requests but through durable interfaces: dashboards, periodic reporting pipelines and application programming interfaces. This shift changes the object of legality. A single request can be assessed by asking whether a stated purpose justifies the particular disclosure. A standing interface, by contrast, creates an ongoing access capability whose intrusiveness accumulates through repetition, aggregation and time. This article develops a three-axis model for encoding the requirement of minimum necessity at the configuration layer of such interfaces. The model treats data fields, extraction frequency and temporal persistence as the minimal auditable surface of standing access, while requiring purpose, recipients, selectors and onward sharing to be recorded in the same authorisation schedule. Using Chinese administrative interface governance as a stress test, especially health-code pipelines and ride-hailing supervisory feeds, the article shows how lawful or plausible access channels can drift through field accretion, cadence escalation and retention extension. EU and US materials are used as design exemplars rather than as complete solutions, illustrating the importance of reasoned requests, strict-necessity limits, duration sensitivity and auditable safeguards. The article concludes by proposing a procedural toolkit: reasoned authorisations, versioned parameter sheets, access logs, renewal triggers, independent audit and remedies for drift. Where feasible, a policy-as-code counterpart can make authorised limits executable at the gateway.

**Keywords:** Government data access; Platform regulation; Data minimisation; Proportionality; Interface governance; China; Auditability

**Online publication:** Jun 29, 2026

## 1. Introduction

Digital platforms have become administrative infrastructures. They organise payments, mobility, communication, labour and consumer transactions, and they concentrate behavioural traces that are

valuable to public authorities. Their legal significance lies not only in private data collection but also in how information flows organise institutional power [1]. The legal problem is no longer only whether a government agency may demand a particular record from a company. Increasingly, the problem is whether the state may rely on a technical channel that makes platform data continuously or routinely available. Once access is engineered as a dashboard, a periodic reporting pipeline or an API, the balance between regulatory capacity and individual rights is shaped by configuration: what is exposed, how often it can be extracted and how long the resulting data remain usable. This is the setting in which traditional necessity review becomes under-specified.

The article argues that standing government access to platform-held data should be reviewed as a configured capability rather than as a sequence of discrete requests. In many legal systems, proportionality and data minimisation are expressed at a high level of abstraction. They ask whether a public purpose is legitimate, whether access is useful or necessary for that purpose, and whether safeguards exist. Those questions remain indispensable. They are, however, too coarse for reusable access channels. A purpose such as transport safety, public health, consumer protection or anti-fraud enforcement can support very different technical arrangements. It may support a weekly aggregate report, a daily batch of pseudonymised records, a live query interface or a full-resolution feed with persistent identifiers. The declared purpose may be constant while the rights intrusion changes dramatically.

To make necessity reviewable, this article proposes a three-axis model: fields, frequency and retention. Fields refer to the schema scope of the interface: identifiers, attributes, behavioural logs, location points, transaction records, risk scores and other derived indicators. Frequency refers to the cadence of extraction: one-off disclosure, scheduled reporting, event-triggered access, polling, streaming or unrestricted querying. Retention refers to temporal persistence: how long extracted data, local copies, linked datasets and access logs are kept, and what deletion or decommissioning triggers apply. These axes are not the whole of legality. Purpose, legal basis, authorised recipients, selectors, population filters and onward sharing are equally important. The claim is narrower: fields, frequency and retention form the minimal configuration surface on which minimum necessity can be specified, implemented and audited.

China is used as a stress test because its regulatory practice makes the administrative turn visible. Platforms are deeply integrated into mobility, payment and communication systems, while supervision often relies on licensing, standard-setting and mandated technical connectivity. Health-code systems during the COVID-19 pandemic and ride-hailing supervision illustrate how emergency or licensing rationales can produce standing access channels that are difficult to contest once they become routine. The point is not that the problem is uniquely Chinese. Rather, China clarifies an infrastructural dynamic that can travel across jurisdictions: once an interface exists, the marginal cost of reuse is low, and expansion may occur through apparently technical changes in schema, cadence or storage windows.

The comparative materials used here are functional rather than exhaustive. EU law supplies examples of reasoned requests, statutory specificity and storage limitation, including the GDPR, the Digital Services Act and the Data Act. CJEU data-retention case law shows how strict necessity is linked to categories of data, persons affected and duration. US Fourth Amendment doctrine, especially Carpenter, shows why time depth and aggregation can change the constitutional character of data access. These materials do not amount to a finished general law of standing state-platform interfaces. They are design exemplars for translating abstract necessity into auditable parameters.

The contribution is therefore practical as well as doctrinal. First, the article recasts minimum necessity as parameterised legality for durable access channels. Second, it identifies a drift grammar: function creep, schema accretion, cadence escalation and retention extension. Third, it proposes a procedural toolkit consisting of reasoned authorisations, versioned parameter sheets, access logs, renewal triggers, independent review and remedies. The aim is not to reject government access to platform data. It is to ensure that access remains tied to a task, bounded by measurable caps and contestable when the interface expands beyond what was authorised.

## 2. From episodic requests to standing interfaces

Government access to platform-held data can be arranged along a spectrum. At one end is case-specific extraction: a subpoena, warrant, inspection order or administrative demand for identified records. In the middle are periodic reporting duties, such as monthly compliance reports or scheduled uploads of operational data. At the other end are standing interfaces: dashboards, APIs, supervisory platforms or near-real-time feeds through which authorities can query or receive data on an ongoing basis. These categories overlap and migrate. A repeated reporting duty may become a stable endpoint; a dashboard may begin as emergency infrastructure and later become a normal administrative tool.

The shift matters because the unit of intrusion changes. Episodic disclosure is usually bounded by the request itself. Its legality can be assessed by asking whether the request is authorised, particularised and proportionate. A standing interface is different. It creates an access capability that can be reused across time and possibly across agencies. Its intrusiveness depends on cumulative features: the breadth of the field list, the number of permitted users, the rate of queries, the volume of exports, retention defaults, linkage identifiers and downstream sharing. Even if each individual query appears modest, repeated access can generate a population-level ledger.

Data minimisation in data protection law helps but does not fully solve the problem. The GDPR requires personal data to be adequate, relevant and limited to what is necessary, and to be kept no longer than necessary [2]. Yet standing access involves a vertical rights relationship between state and individual and a horizontal implementation relationship between state and platform. The platform may be compelled or incentivised to build a channel, authenticate officials, expose fields, retain logs and bear security risk. The affected individuals often cannot see the channel and may not know when their data are accessed. Necessity therefore has to be translated into technical and organisational conditions that operate at the interface.

The same point can be expressed through proportionality. Proportionality review asks whether a measure pursues a legitimate aim, is suitable, is necessary in the sense of having no less restrictive equally effective alternative, and strikes a fair balance. For a durable interface, however, the measure is not simply access in the abstract. The measure is the authorised configuration. A regulator seeking transport safety may need accident statistics, driver licensing status or complaint indicators; it does not necessarily need a continuous route trace for every ride. A public-health authority may need a current risk-status token for venue entry; it does not necessarily need a long-term archive of travel histories. Without configuration-level specificity, proportionality becomes a debate about public purposes rather than about the actual system through which power is exercised. Contextual theories of privacy make the same point from another direction: the normative question is not an isolated data item but whether the flow, role and context of transfer remain appropriate [3].

Interface form also shapes institutional incentives. Pull-based interfaces, such as dashboards and query APIs, enable flexible retrieval but can drift through additional users, broader selectors and larger query volumes. Push-based pipelines, such as scheduled reports, appear more limited but can drift through added fields, shorter reporting windows and longer retention. Emergency systems can be retained as contingencies and then repurposed. Licensing regimes can turn connectivity into a condition of market entry, making refusal practically costly for platforms. In each case, drift may be presented as a technical update rather than a new legal decision.

**Table 1** summarises the basic modalities and the distinct minimality risks associated with each. The table is not a legal taxonomy. It is a way of identifying where the three axes become salient and why a single purpose-level authorisation cannot adequately discipline standing access.

**Table 1.** Modalities of government access to platform-held data

| Modality | Typical mechanism | Minimality risk |
|---|---|---|
| Case-specific extraction | Warrant, subpoena, inspection order or ad hoc administrative request | Overbreadth in the request and weak visibility over downstream reuse |
| Periodic reporting | Batch reports, scheduled uploads or compliance dashboards | Field accretion, shortened reporting windows and silent retention extension |
| Standing interface access | API, supervisory platform, live dashboard or automated alert channel | Continuous monitoring, population-level profiling and procedural difficulty in contesting each access event |

## 3. China as a stress test: Health codes and ride-hailing feeds

Chinese administrative interface governance shows how standing access channels can emerge from different legal and institutional pathways. Two illustrations are especially useful. The first is the health-code infrastructure built during the COVID-19 pandemic. The second is ride-hailing supervision, where platform connectivity operates as a licensing and compliance condition. The two settings differ in purpose, legal trigger and public legitimacy. Yet both show how technical standardisation and administrative convenience make data access expandable along the same three axes.

Health-code systems were deployed as mobility infrastructure. A QR code or similar token indicated whether a person could enter transport hubs, workplaces, residential compounds or public venues. The front end appeared simple: a colour or status signal. Behind it lay data pipelines connecting identity information, travel or location indicators, testing records, quarantine status and local risk rules [4]. During an acute emergency, such systems could be defended as tools for rapid risk sorting. The harder question is what happens after the emergency rationale weakens. Once identifiers, verification practices and back-end data flows are standardised, the interface can be reused for additional administrative purposes.

The Zhengzhou village-bank incident exposed this risk in concrete form. In 2022, authorities reported that officials had improperly assigned red health codes to certain depositors, preventing travel and participation in collective action, and disciplinary measures followed [5]. The episode illustrates how an emergency-born interface can become a behavioural control instrument even without an openly announced change in legal purpose. A system designed for epidemiological risk may be repurposed by changing who can set flags, what predicates trigger status changes or which agencies can query the back end. In three-axis terms, drift occurs when additional linkage identifiers or risk attributes are added, verification becomes continuous or retrospective, and data are retained long enough to support cross-programme use.

A parameterised approach would not ask simply whether health-code data are necessary. It would ask which fields are needed for a defined task. For venue-entry verification, a current status token and minimal validation identifiers may suffice. Full travel histories, testing histories or underlying location traces should require a separate necessity showing. It would ask how often access is needed. Event-based verification at entry points differs from continuous polling or retrospective mass queries. It would ask how long outputs, logs and linked records can persist, and whether deletion is automatic after the public-health trigger ends. Purpose, recipient roles and selectors would be recorded in the same schedule so that a change in authorised agencies or predicate classes becomes a legally salient event.

Ride-hailing supervision shows the same logic outside the emergency setting. China's rules for online car-hailing require platforms to maintain data-processing and exchange capabilities, to connect relevant databases to supervisory platforms and to provide legally authorised regulators with conditions to retrieve and query platform data [6]. Implementing measures for supervision platforms specify operational datasets and timeliness requirements, including high-cadence or real-time transmission for certain categories of data [7]. In regulatory terms, connectivity becomes a condition of market entry and continued operation. In technical terms, it creates a durable data pipe.

Mobility data are especially sensitive because repeated trips can reveal home and work locations, religious practice, health visits, social relationships and political activity. The US Supreme Court recognised in Carpenter that long-term location records may reveal the privacies of life even when individual points seem mundane [8]. Re-identification research also shows why granular mobility traces are difficult to render harmless through simple anonymization [9]. A supervisory authority may need to verify driver licensing, detect safety incidents, monitor service quality or investigate complaints. Those tasks do not automatically justify full route histories, passenger identifiers or open-ended retention.

In ride-hailing supervision, the three-axis model separates regulatory need from surveillance convenience. Field caps distinguish licensing and safety attributes from behavioural reconstruction fields. Frequency caps distinguish periodic statistical reporting from near-real-time feeds capable of continuous tracking. Retention caps tie storage to complaint handling and enforcement windows rather than to indefinite administrative convenience. The model also highlights linkage risk. Stable vehicle, driver, passenger, payment or device identifiers may make it possible to join mobility records with other datasets. Such identifiers should be treated as exceptional fields unless their necessity is specifically justified.

The two illustrations support a broader drift grammar. Function creep occurs when a channel built for one task is reused for another. Schema accretion occurs when additional fields, identifiers or derived scores are added to an interface. Cadence escalation occurs when batch reporting becomes daily, hourly, event-triggered or live. Retention extension occurs when short-term operational data become long-term archives. Linkage creep cuts across all four dynamics by making datasets joinable across programmes and over time. These patterns are not accidental failures. They are predictable effects of standing access, administrative demand and low marginal costs of reuse [10].

## 4. The three-axis model of minimum necessary access

The three-axis model translates the legal idea of minimum necessity into a configuration grammar for standing interfaces. It does not replace legality, proportionality or rights review. It provides the operational surface on which those principles can bite. A lawful authorisation should not merely state that an agency may

access platform data for a public purpose. It should specify the fields that may be exposed, the frequency with which they may be queried or transferred, and the retention rules that govern outputs and logs. Each axis should be measurable, versioned and linked to renewal triggers.

The first axis is fields. Field specification is the closest technical analogue to particularity. It determines which attributes, identifiers, behavioural logs, content fields, metadata and derived indicators are accessible. A well-designed authorisation should contain a data dictionary or schema annex that defines each field, its provenance, its intended use and its sensitivity. It should distinguish raw records from derived indicators; stable identifiers from temporary tokens; and high-resolution fields from aggregate or zone-level outputs. Where aggregated, sampled or privacy-preserving outputs can achieve the regulatory task, they should be the default; differential-privacy techniques are one example of how useful signals may sometimes be released while limiting disclosure risk [11]. Granular location histories, content fields, persistent linkage identifiers and inferential risk scores should require heightened justification.

Fields cannot be separated from selectors. A narrow field list can still enable a dragnet if broad predicate classes are authorised, such as all users in an area, all persons matching a keyword or all accounts linked to a loosely defined risk signal. Conversely, a sensitive field may be acceptable for a narrowly identified case if strict approval and deletion conditions exist. For analytical clarity, selectors are treated as a targeting module in the same authorisation schedule rather than as a fourth axis. The module should specify permitted selectors, prohibited predicate classes, population filters, role-based recipients and onward-sharing pathways. Expanding selectors or recipients should trigger renewed review just as surely as adding a new data field.

The second axis is frequency. Two interfaces with identical fields can differ radically in intrusiveness depending on cadence. A weekly report of aggregate ride volumes is not equivalent to a live feed of vehicle locations. A status check at an entry point is not equivalent to continuous polling of a person's movement record. Frequency is therefore an intensity multiplier. It should be expressed in measurable terms: reporting windows, polling intervals, rate limits, quotas, event triggers and per-role access caps. High-cadence access should be exceptional, tied to a narrow trigger, time-limited and subject to automatic alerts when usage exceeds the authorised profile.

The frequency axis is where quiet expansion often appears first. Query volumes rise; new endpoints are activated; batch reporting becomes near-real-time; emergency flags are used outside the triggering event. Logs should therefore record requester identity, time, endpoint, fields, volume, purpose code and authorisation version for every query or export. The authorised rate limit should be a legal condition, not merely an internal engineering setting. If the gateway is reconfigured to allow more frequent access, the legal schedule should change as well. If the legal schedule is amended, the gateway should be updated and logged. This coupling is the central promise of auditable interface governance.

The third axis is retention. Retention governs how long extracted data, local copies, linked datasets and access logs remain usable. It is central because time changes the character of data power. Long retention permits longitudinal profiles, cross-agency joins and later repurposing. EU data-retention case law has repeatedly linked strict necessity to categories of data, persons affected and duration, rejecting general and indiscriminate retention as a default [12]. In the United States, Carpenter similarly emphasised how historical location records create a detailed chronicle over time [8]. Retention is therefore not an administrative afterthought; it is part of the intrusion itself.

Retention rules should include default windows, deletion triggers, preservation exceptions and

decommissioning conditions. Extracted data should be retained only for the operational task: complaint resolution, safety investigation, audit sampling or emergency response. Longer storage should require a written reason and a renewed authorisation. Logs may need to persist longer than outputs for accountability, but log access should itself be controlled and audited. Emergency-born channels should include an exit plan: when the triggering condition ends, the interface should be disabled, reduced to minimal verification or re-authorised under ordinary standards. Without such rules, temporary infrastructure can harden into a permanent archive.

EU instruments illustrate how parameterisation can be written into law. The Data Act's exceptional-need framework requires reasoned requests that specify purpose, data needed, intended use, duration and, where feasible, erasure expectations [13]. The Digital Services Act creates structured data-access mechanisms for regulators and vetted researchers and recognises that access may be mediated through appropriate interfaces [14]. These regimes are not general codes for all government access. Their importance lies in their design grammar: reasoned access, scope specification, time bounds and traceable procedures. Those features can be generalised to standing state-platform interfaces.

The model can be condensed into a simple test. For any standing access channel, a reviewer should be able to answer five questions. What task justifies the channel? Which fields and selectors are authorised? At what cadence may the channel be used? How long may outputs and logs persist? What happens if any parameter expands? If those questions cannot be answered from the authorisation record and system logs, necessity has not been encoded in a reviewable form.

## 5. Proceduralising necessity: Authorisation, audit and remedies

The three axes become meaningful only if embedded in procedure. A paper statement of necessity does not restrain a standing interface unless it is translated into enforceable settings, audited over time and backed by remedies. The procedural toolkit proposed here has six components: threshold reason-giving, a versioned parameter sheet, technical enforcement, audit logs, renewal triggers and consequences for overreach. Together they turn necessity from an abstract principle into an administrative artefact that can be signed, implemented, sampled and challenged.

First, every standing access channel should begin with a reasoned authorisation. The authorisation should identify the legal basis, the requesting body, the operational task, the less intrusive alternatives considered and the reason those alternatives are insufficient. It should avoid abstract formulae such as public order, market supervision or data security unless those terms are tied to a concrete decision rule. The requester should explain why the task requires platform-held data rather than aggregated statistics, sampling, voluntary reporting, public information or a controlled analysis environment. This threshold explanation creates the record against which later drift can be judged.

Second, the authorisation should include a parameter sheet. **Table 2** provides an illustrative version for a ride-hailing supervision interface. The sheet should be attached to the legal authorisation and should be intelligible to both lawyers and engineers. It should contain a version number, an approved field list, a selector module, per-role recipients, frequency limits, retention windows, logging duties, review triggers and remedies. The goal is not bureaucratic ornament. The goal is to identify the authorised capability with enough precision that a later reviewer can compare what was approved with what the system actually did.

**Table 2.** Illustrative parameter sheet for a ride-hailing supervision interface

| Element | Authorised configuration |
|---|---|
| Purpose and task | Licensing compliance, service-quality monitoring and case-specific safety investigation; no general behavioural profiling. |
| Fields | Driver and vehicle licensing identifiers; trip timestamps; coarse origin-destination zones; complaint and cancellation indicators. Full route traces, persistent passenger identifiers and payment identifiers require separate approval. |
| Selectors and recipients | Named supervisory units and approved analyst roles; prohibited bulk queries by passenger, political event, broad geofence or unrelated agency task without renewed authorisation. |
| Frequency | Routine weekly or bi-weekly batch reporting; case-specific queries subject to per-role quotas and rate limits; real-time feed only under a documented emergency trigger and short sunset. |
| Retention | Operational outputs retained for complaint or enforcement windows, normally 30-90 days; longer retention requires written renewal. Logs retained longer for accountability with strict access controls. |
| Audit and triggers | Every query logged with user, purpose code, fields, volume and schedule version. New fields, higher cadence, added recipients or extended retention trigger independent review. |

Third, authorised limits should be enforced technically where feasible. Closed field lists can be implemented as database views or API scopes. Frequency caps can be implemented through rate limits, quotas and event-trigger gates. Retention can be implemented through automated deletion, irreversible de-linking or expiry flags. Access should be mediated by strong authentication and least-privilege controls. NIST's security and privacy control framework provides useful control families for access management, audit events and configuration change, but generic information-security compliance is not enough [15]. The controls must be tied to the specific parameter sheet.

A policy-as-code counterpart can reduce the gap between legal approval and system configuration. The human-readable parameter sheet can be paired with a machine-readable policy file loaded by the API gateway or access-mediation layer. That file can list approved fields, endpoints, roles, rate limits, quotas and retention commands. It should be version-controlled and cryptographically attestable. When the legal schedule changes, the policy file changes; when the gateway changes, the legal schedule should reflect it. This design does not eliminate discretion, but it makes unauthorised drift easier to detect and harder to deny.

Fourth, audit logs must be usable rather than decorative. A log should record who accessed what, when, through which endpoint, under which purpose code and schedule version, and in what volume. It should also record downstream sharing and deletion where feasible. Platform-side logs and authority-side logs should be reconciled so that missing records, unauthorised endpoints or unexplained query surges become visible. Logs should be tamper-evident and protected from routine users. As high-volume logs can become unread archives, the authorisation should specify a sampling method: periodic review, risk-based sampling, random checks and automatic alerts for anomalies.

Fifth, standing access requires renewal triggers. A one-off authorisation is inadequate because the capability may change without a new public decision. Material changes should include adding a field or linkage identifier, enabling a new endpoint, widening selectors or recipient roles, increasing cadence, exceeding a query budget, extending retention, changing the purpose code or enabling onward sharing. For high-risk channels, these changes should require independent approval. For lower-risk supervisory feeds, they should at least require documented internal review and later external audit. Emergency channels should carry a non-renewal presumption: continuation requires fresh evidence that the trigger persists and, absent new reasons, narrowing rather than expansion.

Sixth, remedies must make parameters enforceable. If an interface exceeds its authorised field list, cadence or retention window, consequences should include suspension or narrowing of the channel, deletion or segregation of unlawfully obtained data, disciplinary or administrative sanctions, and constraints on evidential use where national law allows. Platforms should have safe channels to contest manifestly overbroad demands, insist on renewed authorisation when parameters expand and confidentially escalate concerns to an independent reviewer. Individuals will often lack notice of access, so oversight bodies, courts and auditors must be able to act even where individual challenge is difficult.

Transparency can supplement remedies without exposing sensitive investigations. Authorities can publish aggregate statistics on standing access channels: number of active interfaces, broad purposes, categories of data, renewal decisions, audit findings and instances of suspension or deletion. Such reporting makes drift visible at a system level. It also prevents the vocabulary of necessity from being reduced to private negotiation between government and platform. The Digital Services Act's transparency architecture and audit requirements show how standardised reporting can become part of governance legitimacy, even though the DSA's specific context is platform-risk management rather than general state access.

The toolkit is not a technocratic cure for unlawful or abusive purposes. A precisely configured interface can still be illegitimate if its objective violates fundamental rights or if oversight is captured. Parameterisation also does not address every governance concern, such as discrimination, accuracy, procedural fairness or downstream decision-making. Its value lies in making one crucial dimension of power visible: the configured access capability. Courts and regulators can then ask a concrete question: did the channel operate within the authorised fields, frequency and retention limits, and were expansions justified, approved and remedied?

## 6. Conclusion

Standing access to platform-held data transforms necessity review. When a government agency demands a discrete record, legality can focus on that request. When access is engineered as a reusable interface, the legally salient object is a durable capability. Its intrusiveness is produced cumulatively through repeated use, aggregation and temporal persistence. Necessity must therefore be assessed at the configuration layer. The three-axis model proposed here identifies fields, frequency and retention as the minimal auditable surface of standing access. Field limits determine what is exposed. Frequency limits determine how intensively the channel can be used. Retention limits determine whether data remain available for linkage, profiling and future reuse. Purpose, recipients, selectors and onward sharing must be specified in the same authorisation schedule, but the three axes supply the basic unit of measurable intrusion. China's health-code and ride-hailing examples show why this matters. Emergency tools and licensing pipelines can solidify into standing interfaces. Once the infrastructure exists, drift may occur through field accretion, cadence escalation and retention extension, often without a clear new legal decision. EU and US materials show that rule-of-law systems face the same configuration problem, even if their institutional safeguards differ. Strict necessity, data minimisation, duration sensitivity and reasoned access all point toward the same design requirement: bounded, auditable and renewable access. The practical answer is auditable legality. Authorisations should be reasoned. Parameter sheets should be versioned. Logs should show who accessed what, when and why. Material changes should trigger renewal. Remedies should include suspension, deletion, sanctions and, where appropriate, limits on evidential use. Where feasible, authorised limits should also be expressed as policy-as-

code so that legal schedules and gateway settings remain coupled. The purpose of this architecture is not to eliminate discretion. It is to move discretion from quiet configuration drift to reviewable decisions that can be justified, audited and contested.

## Funding

Tianjin Philosophy and Social Science Planning Project (Project No.: TJFXQN25-05)

## Disclosure statement

The authors declare no conflict of interest.

## References

[1] Cohen J, 2019, Between Truth and Power: The Legal Constructions of Informational Capitalism, Oxford University Press.

[2] Regulation (EU) 2016/679 of the European Parliament and of the Council, General Data Protection Regulation.

[3] Nissenbaum H, 2009, Privacy in Context: Technology, Policy, and the Integrity of Social Life, Stanford University Press.

[4] Zhang X, 2022, Decoding China's COVID-19 Health Code Apps: The Legal Challenges. Healthcare, 10(8): 1479.

[5] PRC CPC News Network, 2022, Zhengzhou Announces Investigation and Accountability for Red Health Code Assignment to Rural Bank Depositors, June 22, 2022.

[6] PRC Ministry of Transport, et al., 2022, Interim Measures for the Administration of Online Car-Hailing Operation Services.

[7] PRC Ministry of Transport, 2022, Measures for the Operation and Management of the Online Ride-Hailing Supervision Information Interaction Platform.

[8] Carpenter v United States, 138 S. Ct. 2206, 2018.

[9] Ohm P, 2010, Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. UCLA Law Review, 2010(57): 1701–1777.

[10] Koops B, 2021, The Concept of Function Creep. Law, Innovation and Technology, 13(1): 29–56.

[11] Dwork C, Roth A, 2014, The Algorithmic Foundations of Differential Privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4): 211–407.

[12] Joined Cases C-511/18, C-512/18 and C-520/18, La Quadrature du Net and Others, EU:C:2020:791.

[13] Regulation (EU) 2023/2854 of the European Parliament and of the Council, Data Act.

[14] Regulation (EU) 2022/2065 of the European Parliament and of the Council, Digital Services Act.

[15] National Institute of Standards and Technology, 2020, Security and Privacy Controls for Information Systems and Organizations, Special Publication 800-53 Revision 5.