

Territorializing Airspace: Reframing Rights–Restrictions–Responsibilities for Low-Altitude Governance

He GAO, Walter Timo DE VRIES

Keywords: Low-altitude Airspace, Territorialization, Rights–Restrictions–Responsibilities, Spatial Justice

Abstract: Conventional land management rests on a two-dimensional logic of parcel, tenure and land use control and, even when extended vertically, remains anchored in land-and-building title: spatial units can be traced to parcels or structures, enabling rights allocation and responsibility assignment. We argue that the emerging low-altitude operating system is reshaping that paradigm. Low-altitude airspace is being organized as a plannable, licensable, operable, and accountable urban space layer. Its governing units are not parcel projections but a time-parameterized, networked operational domain whose availability and constraints are produced through a service layer and whose externalities are negotiated within municipal accountability structures. We ask two questions: (1) How is low-altitude airspace territorialized into a governable layer of urban space? (2) How can a concrete governance model and institutional architecture be constructed on the basis of this layer? Methodologically, we adopt a conceptual-design and synthesis approach, drawing on institutional prototypes across policy corpora, ConOps documents, technical standards, planning guidance, legal doctrines, and evidence on acceptance and equity. We conceptualize territorialization as governability production, and develop a mechanism-levers-enforceability model: the service layer supplies operational mechanisms; cities act through node-based planning and infrastructure levers; and enforceability is secured through traceable responsibility chains, adjudication, and remedies. Within this architecture, the Rights-Restrictions-Responsibilities framework is extended from parcel-based governance to corridors, nodes, time windows, and service interfaces. Spatial justice functions as a design criterion that translates distributive, procedural, and recognitional concerns into executable requirements and triggerable updates. We contribute by (i) theorizing service-layer territorialization as the mechanism that renders low altitude municipal, auditable, and allocable; (ii) proposing an institutional governance architecture linking service interfaces, node-based planning, and enforceability; and (iii) extending RRR to a volumetric, time-parameterized, networked domain, summarized through 8R design criteria.

Territorializing Airspace: Reframing Rights–Restrictions–Responsibilities for Low-Altitude Governance

1 Introduction

Low-altitude urban operations are moving into a phase of routine deployment. The central institutional challenge is no longer flight capability, but how to embed these activities in local governance and sustain an orderly operating regime (EASA, 2021a; ICAO, 2023). As frequency and mission types expand, noise, privacy exposure, safety and environmental risks, and burdens on energy and communications infrastructure shift from occasional impacts to persistent public issues. They also reshape planning and infrastructure decisions through the placement of nodes, the organization of corridors, and the allocation of time windows, with continuing implications for coordination with public safety agencies. At the same time, regulatory authority over airspace is typically concentrated at national or sectoral levels, while cities face more immediate pressure to manage impacts, communicate with the public, coordinate emergency responses, and maintain legitimacy (FAA, 2023; SESAR, 2023). This multi-level mismatch makes questions of spatial justice unavoidable. Benefits and burdens are unevenly distributed, and exposures and risks tend to accumulate near nodes and along corridors. Capacity and permitting rules therefore have clear distributive effects. Without transparent reasoning, traceable responsibility boundaries, and procedures that allow contestation and appeal, social license is unlikely to hold. In this context, access entitlements and priorities, responsibilities for risks and externalities, and restrictions such as no-fly zones, time windows, and capacity limits require an enforceable language of rights, responsibilities, and restrictions. The problem is how to connect Rights-Restrictions-Responsibilities to urban planning instruments, service interfaces, and procedural justice requirements within a volumetric, time-parameterized, networked operating domain.

Research on this issue has developed along three main lines. First, studies on low-altitude traffic management and service-based airspace governance focus on operational rules, digital services, automation, and compliance (EASA, 2021a; FAA, 2023; ICAO, 2023; SESAR, 2023). They examine identity and traceability, rules-as-data, authorization and capacity coordination, conformance monitoring, and auditability, and they increasingly treat governance targets as operational intent, space–time occupancy, and compliance processes rather than aircraft alone (Bauranov & Rakas, 2021; Garrow et al., 2021; Long et al., 2023; Shao et al., 2021). Second, scholarship from an urban planning and infrastructure perspective concentrates on the spatial organization of nodes and corridors and on siting approaches (Cohen et al., 2024; Hijazeen et al., 2025; Mladenović et al., 2024). It emphasizes land-use compatibility, avoidance of sensitive receptors, intermodal connectivity, and requirements for power and communications, and it develops GIS overlays, multi-criteria assessment, and scenario-based modelling to support decisions on node placement and network design (Brunelli et al., 2023; Jin et al., 2024; Lu et al., 2025; Rohrmeier et al., 2025; Schweiger & Preis, 2022).

Third, research on externalities, public acceptance, risk, and legal conflicts addresses the measurement and management of noise, privacy exposure, safety and environmental risks, as well as responsibility allocation and avenues for redress (EASA, 2021b; Orrenius et al., 2025). It also shows that scarcity in capacity and permitting rules has distributive consequences, making fairness, transparency, justification, and appeal procedures central to governance durability (Al Haddad et al., 2020; Bauer et al., 2024; Fried et al., 2024; Yunus et al., 2023).

Despite these advances, several gaps remain. The technical governance and urban planning literatures often run in parallel, with limited explanation of how service layers produce administrable spatial units and how cities, through nodes, permitting, and infrastructure instruments, can incorporate them into local governance. Spatial units are still frequently framed with an implicit two-dimensional land logic, leaving the institutional language for rights, responsibilities, and restrictions in volumetric, time-parameterized, networked domains under-specified. Spatial justice and legitimacy are commonly treated as external evaluations or acceptance concerns, rather than enforceable requirements embedded in authorization, capacity allocation, off-nominal procedures, and audit trails. This makes it necessary to develop an integrative conceptual framework that links service-layer mechanisms, urban governance levers, and conditions of enforceability.

Building on this, the study advances two sequential research questions:

1. How is low-altitude airspace territorialized into a governable layer of urban space?
2. How can a concrete governance model and institutional architecture be constructed on the basis of this urban spatial layer?

Therefore, we develop a conceptual approach that uses spatial justice as a design criterion and Rights-Restrictions-Responsibilities (RRR) as a language for specifying rights, responsibilities, and restrictions. It proceeds in two steps. First, it examines the territorialization of low-altitude operations as a normative process, asking how low altitude is transformed into a governable layer of urban space and how its core operating units—corridors, nodes, and time windows—are shaped by distributive patterns of benefits and burdens, by procedural requirements such as transparency and appeal, and by recognition concerns for affected groups. Second, it uses RRR to organize the governance architecture that follows from this spatial layer. On the rights side, it clarifies conditional access, priority rules, and entitlements linked to node use. On the responsibilities side, it specifies boundaries of obligation across service providers and operators, including conformance monitoring, off-nominal handling, data retention, and public communication. On the restrictions side, it treats no-fly areas, time windows, capacity limits, and buffers around sensitive locations as enforceable constraints, and it shows how cities can operationalize them through node siting and permitting conditions, zoning overlays, public-asset and contract arrangements, and platform interfaces for data exchange and enforcement. The two research questions thus align with these steps: the first addresses the normative logic of territorialization, and the second addresses the construction of a concrete governance model and institutional architecture.

2 Territorialization of Low-altitude Space

2.1 Spatial Units as Governance Objects

Low-altitude operations are better governed as a time-parameterized operational domain rather than as

a fixed spatial container because both occupancy and constraints are inherently dynamic. Aircraft do not 'hold' a bounded volume in a sustained manner; they occupy a sequence of volumetric segments as they move. The intensity of occupancy is shaped by speed, timing, and queuing dynamics, which in turn determines when and where overlaps and conflict risks arise. The scarce governance resource is therefore not a static zone as such, but the admissible flow, risk, and externality load that specific corridor segments and node-associated volumes can accommodate within specific time slices. Permitting, throttling, and allocation in such a setting require time to be embedded as an institutional parameter. Purely geometric controls struggle to address routine conditions such as high-frequency entries, dynamic rerouting, temporary constraints, and congestion spillovers. This is why contemporary low-altitude governance discussions increasingly treat operational ordering as a function of time-based scheduling, capacity allocation, and dynamic constraint publication, maintained through continuous monitoring and structured handling of off-nominal situations (FAA, 2025).

Defining low-altitude space as a 4D networked operational domain reflects an emerging convergence in regulatory and industry thinking. The core idea is to structure traffic and conflict as a routed network, manage availability and priorities through time parameters, and connect rules, authorization, monitoring, and evidence through service-based interfaces. The four concepts used here, corridors, nodes, time windows, and service interfaces, correspond to four necessary components of that representation (SESAR, 2023). Corridors and nodes express the network form that makes routing and flow control possible. Nodes also capture the practical points where cities can institutionalize conditions through land, facilities, and infrastructure. Time windows provide the unit through which capacity and priority can be formalized and updated as conditions change. Service interfaces provide the execution layer that carries rule publication, authorization, situational information, identification, and evidentiary recording in interoperable forms.

Corridors serve as the backbone of the operational domain by turning continuous volumetric occupancy into a manageable structure for movement. A corridor is not merely a geometric passage. It is a conditional set of volumes defined by entry and transit requirements, performance thresholds, separation relationships with sensitive areas, intersection rules, and the way dynamic constraints can be layered onto baseline restrictions. Corridors create governance value in three respects. First, they support publishing geographic constraints in machine-readable forms and combining static and dynamic constraint layers. Second, they provide a clear object for ongoing compliance assessment, enabling determinations of whether operations remain within authorized bounds and when deviations warrant response. Third, they make capacity management locatable. Demand–capacity matching can be applied to corridor segments, conflict points, and time slices, converting throttling and prioritization from abstract principles into executable rules.

Nodes are points of aggregation and distribution within the operational domain, and they are also the points where cities can most readily exercise leverage. Nodes include vertiports, take-off and landing sites, and intermodal transfer points, as well as operational support capabilities such as charging or battery swapping, maintenance, dispatch, and data connectivity. In governance terms, nodes are coupling points for multiple flows: the interface between aerial and ground mobility, the entry point for dispatch and route assignment, and the boundary for operational assurance and safety controls. Nodes often exhibit bottleneck characteristics, with concentrated capacity limits and more predictable conflict patterns, which makes them suitable carriers of institutionalized permitting conditions and operating obligations. Node governance

typically extends beyond the facility footprint to include approach and departure volumes, holding or waiting volumes, and separation buffers, turning the node into a composite object that combines facility, operational volume, and operating rules.

Time windows provide the operational domain with an institutionalized temporal structure, giving capacity management, externality control, and allocation rules an actionable scale. Many governance tensions in low-altitude operations arise less from geometry than from time-varying availability and concurrent occupancy by multiple actors. Time windows therefore function as a basic parameter for capacity governance: operational requests are scheduled in time, quotas and priorities are defined within time slices, and allocations can be updated when constraints change or deviations occur. Time windows also support enforceable externality management. They enable time-specific restrictions for noise-sensitive periods, quiet windows near schools and hospitals, temporary limits around major events, and rapid closures during emergencies. These arrangements extend governance beyond static no-fly rules toward context-sensitive, time-structured constraints.

Service interfaces provide institutional status and an execution pathway for these spatial units, enabling the operational domain to be permitted, monitored, and audited. A recurring feature of UTM and U-space thinking is service-based governance. Identification and traceability, rule publication, authorization and capacity, traffic information, monitoring and response, and logging and evidentiary support are realized through interoperable service functions (ASTM, 2022). Remote identification and tracking capabilities provide a technical baseline for traceability and compliance verification. In this sense, the spatial unit is no longer a geometric form alone (ICAO, 2025). It becomes a composite of geometric boundaries, time parameters, permitting status, and service obligations. This composite is what allows capacity allocation, conflict handling, and accountability to be organized as an executable institutional arrangement.

2.2 The Tripartite Chain and Its Mappings

Organizing low-altitude operations as a governable operational domain requires three elements to hold together within a single institutional structure. First, the service layer provides governance capabilities that allow rules to enter operations in machine-readable forms. Second, cities contribute institutional levers grounded in space and infrastructure, through which the operational domain acquires shapeable boundaries and conditions. Third, a closure arrangement centered on rights, responsibilities, and evidence provides verifiable and correctable support for permitting, throttling, rerouting, response, and redress. Taken together, these elements determine whether low-altitude governance moves from technical feasibility toward institutional credibility.

The first element is the mechanism layer, expressed through the service layer's embedding of rules and procedures into operations. Core capabilities typically include identification and traceability arrangements that establish verifiable links between actors and flight activities; data-based rule expression and constraint publication that translate geographic restrictions, performance thresholds, priority rules, and dynamic constraints into constraint sets readable by operational systems; authorization and capacity management that match operational requests to available capacity and update allocations under congestion, temporary restrictions, or disruptive events; operational monitoring and deviation detection that determine whether activity remains within authorized conditions and trigger responses; off-nominal handling and conflict

management that produce executable response actions under communications loss, deviations, incursions into sensitive areas, or multi-actor conflicts; and logging and audit evidence that fix key decisions, state changes, and response actions into traceable records. The mechanism layer does not replace urban governance. Its role is to provide an executable operational grammar through which spatial and temporal constraints can be published, read, verified, and recorded in a stable manner.

The second element is the lever layer, reflecting the city's capacity to exert structural influence over the operational domain through space, assets, and institutional conditions. Nodes constitute the most direct entry point because they implicate land use, construction permitting, public asset provisioning, intermodal connectivity, and emergency coordination. Cities can institutionalize operational requirements and front-load them into permits and contracts through siting and zoning, development intensity and operating conditions, noise and safety requirements, and the delineation of approach and departure volumes. Corridor-level levers commonly take the form of overlay designations, conditional passage arrangements, avoidance requirements for sensitive spaces, and coordinated provisioning of ground infrastructure as well as communications and energy support, aligned with existing urban spatial-control vocabularies. Time-window levers appear as localized arrangements of operating periods and frequencies—for example, defining operating windows and priorities in response to night-time amenity, quiet needs around schools and hospitals, major events and temporary controls, and emergency-priority channels. Data platforms and public notification arrangements form an additional class of levers by integrating rule explanation, public communication, complaint intake, and incident coordination into urban governance workflows, thereby giving operational-domain governance visibility and communicability.

The third element is the enforceability layer. Its core task is to bind rights, responsibilities, and restrictions to verifiable institutional boundaries and to form a closure loop for evidence and remedies. Minimum conditions include identifiable responsible parties, traceable key decisions, recorded critical state changes, reviewable response processes, the ability for disputes to enter appeal and reconsideration procedures, and remedial actions that can be triggered institutionally and leave verifiable outcomes. Because low-altitude operations rely heavily on automation and platform-based coordination, routine governance tends to manifest as automated execution under predefined rules, while risks and conflicts concentrate in off-nominal and boundary cases. Enforceability provides a clear pathway for attribution and correction in those situations, so that system outputs can be explained, claims can be admitted, and adjudication can rest on evidence.

2.3 RRR Syntax for The Tripartite Chain

Rights–Restrictions–Responsibilities supplies a unified grammar for the tripartite chain and anchors it to 4D units and service interfaces. The mapping can be expressed through three sets of questions:

Rights: Who may enter which corridors and nodes, within which time windows, and with what priority or quota entitlements? Who may access which service interfaces and obtain which categories of operational and constraint information? How are entry and operating conditions institutionalized within permits and operating terms?

Responsibilities: What obligations fall on operators, platform service providers, node operators, and relevant city agencies in relation to noise, privacy, safety, data governance, and emergency coordination?

How is performance of these obligations recorded, and which data and logs constitute admissible evidence? Who triggers, executes, and bears the consequences of off-nominal handling and redress?

Restrictions: Which geographic restrictions, thresholds, buffers, and altitude-layer requirements apply to corridors and nodes? Which capacity ceilings, time-of-day limits, frequency controls, and priority rules apply to time windows? Which equipment requirements, algorithmic rules, and dynamic constraints must be executed through interfaces and remain updatable?

The value of this mapping lies in decomposing institutional questions into implementable objects, parameters, and evidentiary requirements. Corridors and nodes carry spatial boundaries and passage conditions; time windows provide the temporal scale for capacity governance and distributive ordering; service interfaces carry status, rules, records, and audit functions. In this way, service-layer mechanisms, urban levers, and accountability closure can be written into a single institutional architecture, with verifiability, attributable responsibility, reviewability, and redress serving as consistent criteria for enforceable governance.

3 Governing Architecture: Service Layer, Node-Based Urban Levers, and Enforceable Closure

3.1 Minimal Functional Set and Institutional Role

Service-layer mechanisms sit at the operational core of low-altitude governance because they turn rules, constraints, and procedures into interoperable service capabilities that run continuously. This shifts governance from textual requirements and ad hoc coordination into routine execution. The institutional role can be summarized in three tasks. First, constraints must be expressed as machine-readable objects with coherent versioning, so that “what rules applied” is unambiguous during operations and in later review. Second, entry and passage must be operationalized as allocable and adjustable authorization states that can be scheduled, throttled, or withdrawn under changing risk and congestion. Third, monitoring, response, and records must be connected into a recoverable chain, so that key decisions and state changes rest on verifiable evidence. In this sense, the service layer moves low-altitude space from merely flyable to permit-able, throttle-able, monitor-able, and accountable.

A minimal functional set has converged around six capabilities: identification and traceability; constraint publication and version management; authorization and capacity management; traffic information, monitoring, and deviation detection; off-nominal handling and conflict management; and logging and audit support. Identification and traceability provide the object basis for governance. Routine operations require that activities can be linked to identifiable parties and vehicles, enabling authorization, enforcement, and responsibility attribution. The point is not a specific technical solution but a stable information minimum and consistent exchange practice: actor and vehicle identifiers, availability of position and trajectory information at appropriate granularity, and defined boundaries for sharing when privacy or commercial sensitivity is implicated. Without this layer, authorization and accountability risk collapsing into unverifiable claims.

Constraint publication and version management determine whether rules can be executed at scale. Low-altitude constraints are strongly time-varying and context-dependent. Fixed prohibitions are only a baseline; operational order depends on temporary restrictions, dynamic closures, event windows, degraded-mode operations, and congestion- or risk-triggered constraints. Making these constraints operational requires expressing them as computable sets with explicit validity periods, geographic scope, update frequency, and

quality control. Version consistency is especially important because it provides the shared reference needed for compliance verification and ex post reconstruction. When a denial, reroute, or suspension is contested, the first question is not only what the rule said in abstract, but which constraint set was in force, where, and when.

Authorization and capacity management transform access into allocable entry rights. Capacity is rarely experienced as an aggregate ceiling bounded by a static airspace perimeter. It binds at bottleneck corridor segments, node-associated volumes, and time slices defined by admissible flow and risk thresholds. The service layer maps requests into concrete space–time scopes, aligns authorization states with capacity constraints, and turns entry into an operational status that can be scheduled, adjusted, and—when necessary—withdrawn. Authorization therefore has distributive implications by design: who may enter, when, under which conditions, and who receives priority under congestion or temporary restrictions. The institutional significance is that allocation becomes a recurring decision stream rather than a one-off permit.

Traffic information, operational monitoring, and deviation detection sustain the ongoing effectiveness of constraints and authorizations. For restrictions to constitute order, the system requires continuous state information and interpretable deviation assessment: whether an operation remains within permitted bounds, whether alerts should be triggered as boundaries are approached, and whether deviations require escalation. The technical emphasis is not comprehensive surveillance but a minimal monitoring loop: continuity of state data, deviation thresholds that can be interpreted and reviewed, traceable alert triggers, and a clear linkage from detection to response actions.

Off-nominal handling and conflict management determine whether governance can contain public risks at boundaries. High-frequency, multi-actor operations concentrate pressure in recurring scenarios: deviations in route, time, altitude, or performance; loss or degradation of communications, navigation, or positioning; incursions into sensitive areas; convergence on conflict points producing spatial or capacity conflicts; failures of critical service components undermining authorization integrity; and emergencies requiring temporary closures and priority channels. These scenarios share acute time sensitivity. Effective handling must be organized as executable procedure rather than ad hoc coordination. At minimum, procedure requires four elements: trigger conditions, decision authority, action boundaries, and restoration conditions. Trigger conditions specify when intervention is mandatory, using thresholds for deviation magnitude, time-to-boundary, safety margins, conflict likelihood, or risk levels. Decision authority specifies who may issue directives and how contradictory commands are avoided in multi-actor coordination. Action boundaries define permissible interventions such as rerouting, holding, speed adjustments, entry restrictions, landing requirements, temporary closures, or transitions into emergency modes. Restoration conditions define when and how normal operations resume, including rules for lifting restrictions, recalculating authorizations, re-verification requirements, and compensatory arrangements where appropriate. This proceduralization shifts exception handling from experience-based reaction to institutionalized response, which is a prerequisite for routine governance.

Logging and audit support provide the practical basis for accountability and review. Accountable governance depends on recoverable records of key decisions, critical state changes, and significant response actions, enabling oversight checks, dispute review, and remedy processes to proceed on shared evidence. What must be preserved is not only outcomes, but the applicable constraint version, capacity state, ordering

logic, trigger thresholds, directive pathways, and traces of human intervention. Record quality shapes institutional credibility: without records, allocation and response decisions appear as opaque system outputs; with records, explanation, review, and correction become actionable.

Finally, the boundary of the service layer matters. It does not replace urban planning instruments or public value choices. Decisions about node siting, buffers and thresholds around sensitive places, operating windows, and priority orientations remain matters to be defined through urban levers. The service layer translates these requirements into executable constraints and produces verifiable, reviewable records so that governance moves from declaration to stable execution.

3.2 Urban Levers via Nodes

Nodes are the city's primary entry point into low-altitude governance because they combine network structure with institutional structure. In network terms, nodes determine where flows converge, disperse, and become bottlenecks, thereby shaping the backbone of corridors and the geography of conflict points. In institutional terms, nodes are where conditions can be stabilized, inspected, and updated: cities can attach executable requirements to scale, operating hours, supporting facilities, information duties, and emergency readiness, and link these conditions to service-layer execution. Node configurations and node conditions therefore reshape both accessibility and exposure patterns across the city.

Treating nodes as simple take-off and landing sites understates their governance significance. Operationally, a node comprises a set of volumes and interfaces that are directly tied to flight processes: approach and departure volumes, separation and avoidance requirements near take-off and landing, and holding volumes when queuing is necessary. It also includes organizational interfaces that connect aerial operations to ground transfers, passenger or cargo flows, maintenance, and support services. In addition, nodes depend on ground-side prerequisites—energy provision, communications and data connectivity, on-site safety arrangements, and emergency response capacity—which jointly determine effective throughput and acceptable risk. These dependencies create stable “handles” for cities: by governing node conditions, cities indirectly govern the operational feasibility and intensity of low-altitude traffic.

Cities are able to act on nodes because node decisions can be carried through existing planning and regulatory vehicles. Location and form are shaped through land-use planning and construction permitting. Scale and throughput can be constrained through development and operating conditions. Temporal patterns can be shaped through opening hours and operating windows. Accessibility and service populations can be influenced through ground-transport integration and public-asset provision. Controllability under disruption can be strengthened through requirements for emergency access, on-site assurance, and coordination arrangements. Embedded in permits, operating agreements, and public-facility conditions, these variables become inspectable, enforceable, and adjustable over time. This is how cities can exert structural influence even when they do not hold comprehensive airspace authority.

Nodes also concentrate public-interest issues and, in turn, reshape corridor networks. Disturbance and exposure often accumulate around nodes and become visible as disputes, including noise, safety concerns, and privacy sensitivities. Corridor connectivity is largely derived from node distribution, throughput, and the constraints implied by approach/departure volumes and conflict density. Concentrated nodes tend to push capacity pressure onto a small set of bottleneck segments; dispersed nodes broaden coverage but increase the

complexity of avoidance design and time-window governance. For these reasons, nodes are best treated as composite governance objects: the facility itself, the operational volumes tied to it, and the operating conditions that can be institutionalized through permits and operating terms.

Corridors and time windows are operational structures derived from node distribution, avoidance requirements, and bottleneck capacity. Connectivity demands between nodes orient the main network directions; avoidance of sensitive receptors constrains feasible geometry; and approach/departure volumes and conflict-point density around nodes determine where capacity constraints concentrate. As a result, corridor skeletons carry an inherently urban character: they are organized around urban functions while redistributing accessibility and disturbance exposure within the city. Urban shaping is therefore not “drawing a line,” but representing corridors as manageable sets of operational volumes aligned with surface land uses and receptor sensitivity. This includes buffers and graded restriction intensity around sensitive receptors; stricter altitude layers or passage conditions near hospitals, schools, and heritage districts; and more stable passage bands over areas with higher tolerance, such as transport corridors, industrial zones, and logistics districts. Differentiation across altitude layers is particularly important because the same horizontal location can have distinct operational availability at different heights.

Time windows provide a parallel governance pathway by embedding operational order into the city's rhythms and public needs. Night-time amenity, sensitive periods for schools and hospitals, major events, temporary controls, and emergency-priority passage all call for actionable time-parameterized arrangements that can be written into permitting conditions. Time windows enable context-sensitive adjustment of externalities and risk without rewriting long-term spatial structure. They also extend capacity governance beyond spatial availability into time-of-day and frequency-based availability, supporting finer alignment with public requirements.

For these spatial and temporal expressions to become operational order, they must be parameterizable in forms that the service layer can publish and execute. Buffers require explicit distances or thresholds; restriction intensity must map to checkable condition sets; time windows must be defined as computable validity intervals with priority rules; and corridor segments and node-associated volumes must have traceable extents and versions. Only when planning terms can be translated into constraint sets do corridors and time windows remain consistent under dynamic restrictions, congestion recalculation, and off-nominal handling.

Cities' most reusable toolkit follows from this logic: public requirements should be embedded into institutional vehicles that can be executed in day-to-day operations. Three vehicles are especially stable. First, permit conditions can front-load node constraints: caps on scale and throughput, operating windows, approach/departure boundaries and directional requirements, and buffering around sensitive places, expressed in inspectable terms. Second, infrastructure and operating conditions turn dependencies into controllable boundaries: site access, charging capacity, communications reliability, on-site safety provisions, and emergency assurance affect throughput stability and response capacity, and can be tied to compliance and performance obligations. Third, information and procedural conditions institutionalize explanation and review: coherent public notices for rule changes and temporary restrictions, accessible complaint intake, and workable pathways for reason-giving and review when denials, throttling, or rerouting occur.

A Spatial Decision Support System (SDSS) can integrate these elements as an organizing framework for planning and permitting decisions. Building on established GIS-based decision-support traditions, SDSS

links data, indicators, scenario comparison, and traceable records. In this paper, SDSS is adapted in two ways: the decision object is expanded to the 4D operational domain of nodes, corridors, and time windows; and the output is upgraded from plan ranking to a clause-based package that can be incorporated directly into permits and operating arrangements. The package structures spatial definitions (node volumes, corridor boundaries, impact extents), temporal rules (window rosters and validity conditions), operational settings (eligibility and capacity arrangements), and governance requirements (duties, records, and review hooks). Where useful, an RRR-structured output format can keep eligibility conditions, restriction parameters, and duty requirements in a consistent grammar, so that spatial choice and institutional implementation sit on the same decision chain. Reviewability then targets not only why a node is selected, but why particular thresholds, windows, and obligations are adopted—supported by minimal records on data definitions, parameter-setting choices, and clause-generation linkages that make disputes contestable and corrections implementable.

3.3 Enforceable Closure

Noise, privacy exposure, and safety risks in low-altitude operations are spatially clustered and time-sensitive. When governance relies on ex post determinations and ad hoc coordination, expectations become unstable and enforcement becomes uneven. This section treats externality governance as an enforceability task: defining impact boundaries through monitorable indicators, specifying triggers through thresholds and time windows, stabilizing response pathways through graded procedures, and keeping all elements consistent with the operational arrangements of nodes, corridors, and time windows. Operationally, externalities can be grouped into five categories: noise and vibration disturbance; privacy and data exposure; ground safety and fall risk; environmental and energy loads; and communications and data-service pressures. Effective parameterization follows three boundary types. A spatial boundary is expressed through buffers, graded intensity zones around sensitive receptors, directional requirements, and altitude-layer conditions. A temporal boundary is expressed through quiet-hour windows, sensitive-period limits, event windows, and temporary control regimes. An intensity boundary is expressed through frequency caps, capacity ceilings, and operating-condition thresholds that support throttling and constraint escalation. Privacy and data exposure typically require combined place–mission rules, such as limits on recording near sensitive sites, requirements on data collection and retention for specific data types, and constraints on interface access. Ground-safety grading depends on the combination of consequence severity and response conditions, including proximity to critical facilities, population exposure, and emergency reachability, which in turn shape passage conditions and time-window settings.

Thresholds and time windows are only enforceable when paired with graded response rules and explicit trigger and restoration conditions. Responses can be organized as escalating layers by scope and urgency: alerts as thresholds are approached; throttling and tighter operating conditions under persistent proximity or localized congestion; rerouting and avoidance as conflict likelihood rises or sensitive boundaries are approached; suspension or temporary closure under service degradation, loss of link, or accumulated high risk; and area closures during major events or emergencies. Each layer requires computable or verifiable triggers and observable restoration conditions, so that temporary measures do not drift into arbitrary extensions and operators can anticipate how normal operations resume. Enforceability also depends on monitoring baselines and minimal evidence to support triggering, review, and later correction. A minimal

record set includes the applicable constraint version and key state snapshots; timestamps, directive pathways, and execution outcomes for alerts and response actions; reasons and outcomes for human intervention; structured records of noise-, exposure-, and privacy-related events; and post-event reports for off-nominal incidents. Records ensure that disputes can be assessed against shared facts and that subsequent adjustments can be anchored in specific parameters, triggers, and response steps. Remedy arrangements should align with this structure so externality governance does not reduce to restrictions alone. Rapid mitigation addresses high-frequency disturbances through temporary throttling, short-term rerouting, tightened operating hours, or enhanced on-site assurance. Institutional correction addresses persistent exposure through recalibrated buffers and intensity gradations, revised time windows and capacity ceilings, and modified node operating conditions and priority regimes. Where impacts are sustained and attributable, governance also requires an intake channel, explicit timelines, and executable compensation or substitute-service arrangements.

Low-altitude operations are routinely organized through multi-actor coordination. Operators execute missions; service providers publish constraints and execute authorization and capacity; node operators maintain facilities and associated volumes; city agencies embed public requirements through permits, communication, and emergency coordination; and higher-level regulators provide overarching frameworks and supervision. Enforceability requires shifting responsibility from *ex post* blame to *ex ante* structural allocation so that key events can be reconstructed as continuous sequences. Responsibility boundaries are best specified by function and tied to verifiable duties: operators for operational compliance and required information submission; service providers for constraint publication and version consistency, authorization and capacity execution, monitoring integrity, and record retention; node operators for facility and ground-order maintenance, access and support capacity, and emergency readiness; city agencies for permit conditions, public-facing notices, complaint intake and review procedures, and inter-agency coordination; and regulators for minimum compliance frameworks and supervision of service-provider performance. An attributable chain should be built around traceable responsibility nodes across the process: constraint updates and entry into force; authorization and capacity allocation; compliance monitoring and deviation assessment; off-nominal handling and restoration; and incident reporting and outward explanation. For each node, four elements should be specified: triggering authority, execution duty, consequence-bearing, and the actor responsible for outward explanation. Evidentiary chains provide the operational basis for attributable responsibility and should be configured to match these nodes: constraint version and scope; authorization status and capacity state; threshold-trigger records; response logs with timestamps and directive pathways; and structured fields for incident reconstruction. Oversight can then focus on version consistency, execution consistency, and off-nominal completeness, feeding graded corrective measures — degraded-mode operation, suspension, restoration, or condition updates—back into permits and operating terms.

Conflict adjudication provides the procedural interface that turns externality claims and operational disputes into reviewable case units, and connects outcomes to parameter updates and remedies. Disputes cluster around exposure and disturbance on the public side, and around eligibility, throttling, and operating conditions on the operational side. The institutional task is to provide pathways that are evidence-based and triggerable: intake, evidentiary assembly, decision, layered review, and remedy. Intake channels can be differentiated between public-facing claims (disturbance, exposure, risk) and operations-facing disputes (denials, throttling, rerouting, suspensions). Time limits should track urgency: fast decisions for temporary

restrictions paired with an ex post window for explanation and review. Evidentiary expectations should be feasible and verifiable: monitoring outputs and exposure baselines for public claims; authorization status, constraint version, capacity state, trigger records, and response logs for operational disputes. Reason-giving should cover three items: applicable rules and parameter scope; trigger basis and evidence; and the rationale for restriction intensity, duration, and exit conditions. Layered review reduces over-concentration of rule-making, execution, and interpretation: a first layer provides rapid correction by the implementing actor; a second, more independent layer examines evidence chains and rule application and can require parameter recalibration, clause updates, or remedies; high-impact disputes retain an external review backstop. Remedy instruments should correspond to conflict types and take the form of adjustable operating conditions: buffers, quiet-hour windows, frequency and capacity ceilings, directional and altitude-layer constraints; data boundaries and interface permissions; priority rules and quota structures with correction pathways; and, where necessary, time-limited restrictions with explicit restoration conditions.

Justice and legitimacy provide the design check that turns thresholds, responsibility chains, and adjudication procedures into a sustainable updating capacity. Spatial justice functions here as an enforceable criterion, supplying verifiable diagnostics for rule-setting, rule execution, and rule revision, and linking these diagnostics to correction triggers. Distributive justice requires comparable baselines for coverage and exposure and periodic evaluation of cumulative exposure around nodes, exposure boundaries for sensitive receptors, and spatial disparities in service access. Triggers for rebalancing include persistent threshold exceedance, trends toward concentrated exposure, and misalignment between coverage gains and burden concentration; corrections take the form of parameter recalibration—buffers, intensity gradations, time-window arrangements, capacity ceilings, and revised node operating conditions. Procedural justice requires reason-giving and review entry points for key decisions, supported by reconstructable records; temporary measures can be fast but must remain reviewable ex post, and high-impact matters should route to a relatively independent review unit. Recognition justice requires explicit protection for sensitive receptors and critical public functions, together with procedural accessibility so that complaint and review channels remain usable under information asymmetries. These criteria must be embedded in an update loop structured around periodic evaluation and event-driven review, supported by version management. Monitoring outputs, complaint statistics, and post-incident reviews feed a revision workflow; outputs take the form of parameter updates and clause revisions with minimal change justification. Under this structure, enforceability is not reduced to rules on paper: it is secured through identifiable responsible parties, reconstructable evidence chains, accessible review pathways, and triggerable remedies and updates.

4 Conclusion: Integrated Framework

As a conclusion, the conceptual argument must be consolidated into reusable institutional outputs. This article conceptualizes low-altitude airspace as an urban spatial layer organized and territorialized through a service layer, with governance centered on a 4D operational domain composed of nodes, corridors, and time windows, and with enforceable boundaries and procedures formed through planning and permitting. On this basis, the chapter distills a closing closed-loop governance architecture to clarify how service-layer mechanisms, urban levers, responsibility chains, and adjudication procedures operate coherently within a single institutional sequence. It then advances an 8R set of design principles—Retraceable, Robust,

Recognizable, Responsive, Resilient, Respected, Reflexive, and Reliable—embedding these requirements in key institutional sites such as rule expression, capacity allocation, record-keeping and review, redress, and iterative updating, and deriving from them a set of testable directions for future research.

The integrated framework advanced in this study can be stated as a stable closed-loop chain. The service layer renders the low-altitude operational domain governable; the 4D operational units carry boundaries and states; cities shape corridors and time windows through nodes and their institutionalized conditions; responsibility chains and adjudication procedures turn rule execution into attributable and remediable processes; and spatial justice functions as a continuous calibration standard that brings distributive outcomes and procedural performance into an updating mechanism. The key claim is that governance does not depend on a static partition of “airspace itself,” but on the dynamic organization of operational availability. Rules enter authorization and capacity arrangements in computable form; operational states are published and verified through interfaces; and disputes and off-nominal situations are routed through review and correction procedures—maintaining a consistent institutional sequence across planning, permitting, operations, and accountability.

The framework comprises five nested layers, with closure achieved through ‘interfaces—versioning—records’. The first layer is the service layer’s capacity to publish rules and states: constraints, authorizations, capacity conditions, and off-nominal handling requirements are made available in data form, together with traceable version histories. The second layer is the spatial-object layer of the 4D operational domain: node volumes, corridor segments, and time windows serve as the minimum governance units, carrying applicability scope, intensity parameters, and availability status. The third layer is the urban shaping layer: through institutional vehicles such as node permit conditions, ground-side assurance, and contractual clauses, cities embed public requirements into nodes and their associated volumes, thereby indirectly structuring corridor organization and temporal ordering. The fourth layer is the enforceability layer: responsibility allocation, evidentiary chains, and adjudication-and-redress procedures ensure that triggering, response, explanation, and remedy can be reconstructed and admitted to review. The fifth layer is the justice-calibration layer: distributive, procedural, and recognitional standards define evaluation baselines and triggers, translate inputs from monitoring, complaints, and post-incident reviews into parameter recalibration and clause updates, and—via version management—write these revisions back into the first layer. In this way, low-altitude governance is organized as an updatable, auditable, and correctable institutional cycle rather than as a one-off plan or a one-directional technical control.

4.1 Implementation Sites Based on 8R Criteria

The 8R criteria compress the closed-loop governance architecture proposed into a deployable set of design requirements (Gao et al., 2026). Each ‘R’ maps to a specific institutional site and execution object, and can be operationalized across key points such as rule expression, capacity allocation, record-keeping and review, redress, and iterative updating—providing a practical basis for testing and calibrating enforceability and legitimacy in low-altitude governance.

Retraceable: Critical decisions and critical response actions should be reconstructable as a continuous sequence, at minimum linking the applicable rule version, authorization and capacity states, triggering basis, response actions, and outcomes. Retracement does not require exhaustive data capture; it requires a minimal

set of reconstructable records at key nodes so that review, audit, and correction can proceed on a shared factual basis.

Implementation sites: service-layer logs and audit interfaces; constraint version management; required response-log fields and retention periods.

Robust: Rules and procedures should retain a consistent execution logic under stress conditions such as congestion, deviations, and temporary controls, with clearly specified pathways for degraded operation and restoration. Robustness emphasizes predictability at boundary conditions, preventing loss of control at critical moments or inconsistent execution across actors.

Implementation sites: triggers and exit conditions for dynamic constraints; conservative capacity settings for bottlenecks; thresholds for degraded-mode operation and restoration.

Recognizable: Responsible parties, decision-makers, and reviewers should be externally identifiable and internally locatable, with clear boundaries for triggering authority, execution duties, consequence-bearing, and explanatory responsibility. Recognizability reduces the space for “responsibility drift” in dispute handling and stabilizes public communication and cross-actor coordination.

Implementation sites: ex ante allocation of duties and responsibility nodes; clearly designated actors for intake and explanation; public-facing disclosure of channels and points of contact.

Responsive: The governance system should be capable of timely intervention in time-sensitive situations and should complete a response loop through graded rules, preventing risk and disturbance from accumulating. Responsiveness requires short response chains, explicit triggers, clear command paths, and workable interfaces with urban emergency coordination.

Implementation sites: graded ladder from alerts to throttling, rerouting, suspension/closure; emergency coordination protocols; availability of on-site assurance and communications.

Resilient: The system should maintain essential functions and recover under shocks and disruptions through both network redundancy and institutional redundancy, including alternative corridors, alternative nodes, backup time windows, and operable priority passage for emergencies. Resilience requires institutionalized restoration conditions and re-assessment thresholds, rather than restoration through ad hoc coordination.

Implementation sites: redundant network design and alternative routing; ground-side assurance at critical nodes; explicit restoration conditions and re-evaluation rules.

Respected: Expectations of amenity, privacy, and safety should have an enforceable place in the institutional design and be translatable into boundaries, thresholds, and time-window arrangements. Being “respected” does not imply zero disturbance; it requires operable priority protection for high-sensitivity places and critical public functions, and institutional receptivity to affected parties’ baseline expectations.

Implementation sites: buffers and graded zones for sensitive receptors; quiet-hour windows; boundaries for data collection and retention; priority conditions for public-interest missions.

Reflexive: Rules and parameters should be able to enter an updating process after monitoring, complaints, review outcomes, and post-incident learning, with preserved rationales and version histories. Reflexivity requires triggerable, traceable, and explainable updates—translating external pressures into parameter recalibration and clause revision rather than relying on case-by-case handling alone.

Implementation sites: dual triggers of periodic evaluation and event-driven review; version governance;

minimal change rationales and write-back mechanisms.

Reliable: Rule publication, authorization execution, information disclosure, and procedural operation should remain consistent and predictable across actors and contexts, preventing “drift” that erodes trust. Reliability emphasizes stable institutional operation: standardized rule expression, consistent interface information, and predictable timelines and pathways so that participants can form durable expectations and compliance behaviors.

Implementation sites: data-based rule expression and consistent publication; standardized clause packages; stable review timelines and procedures.

4.2 Implications and Research Agenda

The contribution lies in providing a reusable institutional language that aligns planning, permitting, operations, and accountability around the same spatial objects and decision chain. Rather than treating low-altitude governance as a set of parallel topics—technical systems, siting tools, and externality disputes—the paper locates governability in an enforceable institutional organization: rules enter authorization and capacity arrangements in computable form; node conditions constitute a stable entry point for urban intervention; responsibility chains and adjudication procedures route disputes into reviewable tracks; and justice criteria are internalized through version governance and parameter updating so that distributive outcomes and procedural performance become governable inputs. On this basis, future research can develop testable propositions across three directions — spatial distribution, procedural performance, and off-nominal handling—and translate them into observable indicators and comparative designs.

Proposition 1: Node–Time-Window Configuration and Distributive Outcomes. Holding demand scale and fleet capability constant, more concentrated nodes with looser time windows improve coverage efficiency but are more likely to intensify spatial concentration of exposure. More dispersed nodes with tighter time windows reduce exposure concentration but may increase operating costs and time uncertainty. Related research can construct joint metrics across coverage, exposure, and cost; compare distributive outcomes under different combinations of node layouts and time-window rules; and identify receptor-specific thresholds and rebalancing triggers around sensitive sites.

Proposition 2: Capacity Allocation Rules and Legitimacy Pressure. Under bottlenecks and throttling conditions, capacity allocation rules that incorporate reason-giving, accessible appeals, and reviewable procedures reduce complaint intensity and dispute spillover while improving rule stability. By contrast, ad hoc coordination or opaque priority ordering is more likely to generate persistent contention. Related research can typologize allocation mechanisms and operationalize legitimacy pressure using indicators such as complaint rates, review-trigger rates, frequency of temporary controls, and the rate of rule changes, and then estimate the marginal effects of procedural elements such as reason-giving, review layering, and timeline commitments across different mechanisms.

Proposition 3: Retraceable Responsibility Chains and Off-Nominal Performance. Architectures with clearer responsibility boundaries and more reconstructable evidentiary chains reduce response and restoration times in off-nominal events, lower secondary risks, and shorten post-incident dispute duration and redress costs. Conversely, broken evidentiary chains or ambiguous responsibility nodes amplify risk spillover and institutional distrust. Related research can treat the trigger–response–restoration sequence as the unit of

analysis, develop measures for responsibility-chain clarity and evidentiary-chain completeness, and evaluate their effects on response latency, restoration duration, repeat-event rates, and dispute-handling costs.

Proposition 4: Justice-Calibration Capacity and Policy Sustainability. Institutionalizing distributive, procedural, and recognitional justice through periodic evaluation and event-driven review, supported by version governance and minimal change justification, increases policy sustainability and public acceptance and reduces policy oscillation driven by accumulated case-level frictions. Related research can compare rule stability, public trust, and dispute density under different strengths of updating institutions, and identify which monitoring baselines, review statistics, and post-incident findings most effectively translate into parameter recalibration and clause revision.

In sum, the closed-loop architecture and the 8R design principles offer an enforceable institutional articulation for low-altitude governance and provide clear objects, variables, and comparative pathways for subsequent research. Future work can test these propositions across different institutional prototypes and urban contexts, while further refining operable templates for node governance clauses, capacity rules, and evidentiary standards—thereby advancing the conceptual framework into an extensible repertoire of urban governance design options.

References

- Al Haddad, C., Chaniotakis, E., Straubinger, A., Plötner, K., & Antoniou, C. (2020). Factors affecting the adoption and use of urban air mobility. *Transportation Research Part A: Policy and Practice*, 132, 696–712.
- ASTM. (2022, July 13). *Standard Specification for Remote ID and Tracking*.
- Bauer, M. W., Bartels, S., & Redmann, D. (2024). *UAM Community Noise Impact Studies in Project PAULA*.
- Bauranov, A., & Rakas, J. (2021). Designing airspace for urban air mobility: A review of concepts and approaches. *Progress in Aerospace Sciences*, 125, 100726.
- Brunelli, M., Ditta, C. C., & Postorino, M. N. (2023). New infrastructures for Urban Air Mobility systems: A systematic review on vertiport location and capacity. *Journal of Air Transport Management*, 112, 102460.
- Cohen, A., Shaheen, S. A., & Wulff, Y. (2024). *Planning for Advanced Air Mobility*. American Planning Association.
- EASA. (2021a, April 23). *Easy Access Rules for U-space (Regulation (EU) 2021/664)—May 2024*.
- EASA, N. (2021b). Study on the societal acceptance of urban air mobility in Europe. *European Union Aviation Safety Agency and McKinsey & Company*.
- FAA. (2023, August 31). *Urban Air Mobility (UAM) Concept of Operations | Federal Aviation Administration*.
- FAA. (2025, May 2). *Unmanned Aircraft System Traffic Management (UTM) | Federal Aviation Administration*.
- Fried, T., Goodchild, A., Browne, M., & Sanchez-Diaz, I. (2024). Seeking equity and justice in urban freight: Where to look? *Transport Reviews*, 44(1), 191–212.
- Gao, H., de Vries, W. T., Shen, W., & Qu, W. (2026). From sustainable to responsible: A biopolitics reinterpretation of ‘8 R’ model. *Land Use Policy*, 161, 107875.

- Garrow, L. A., German, B. J., & Leonard, C. E. (2021). Urban air mobility: A comprehensive review and comparative analysis with autonomous and electric ground transportation for informing future research. *Transportation Research Part C: Emerging Technologies*, 132, 103377.
- Hijazeen, S., King-Yin Cheung, T., Lei, Z., & Hayward, J. A. (2025). Integrating vertiports into Australian airports—A comparative literature review of regulatory frameworks from CASA, FAA, and EASA. *Transport Policy*, 172, 103754.
- ICAO. (2023). *Unmanned Aircraft Systems Traffic Management: UTM (Edition 4)*. CRC Press.
- ICAO. (2025). *Remote ID Number Registration*.
- Jin, Z., Ng, K. K. H., & Zhang, C. (2024). Robust optimisation for vertiport location problem considering travel mode choice behaviour in urban air mobility systems. *Journal of the Air Transport Research Society*, 2, 100006.
- Long, Q., Ma, J., Jiang, F., & Webster, C. J. (2023). Demand analysis in urban air mobility: A literature review. *Journal of Air Transport Management*, 112, 102436.
- Lu, Y., Zeng, W., Wei, W., Wu, W., Jiang, H., Lu, Y., Zeng, W., Wei, W., Wu, W., & Jiang, H. (2025). Vertiport Location Selection and Optimization for Urban Air Mobility in Complex Urban Scenes. *Aerospace*, 12(8).
- Mladenović, M., Niemi, L., Saif, A., & Honkavaara, E. (2024). *Development of a geospatial decision-support tool for Urban Air Mobility landing and launch site location planning: Analysis, framework and technical setup*.
- Orrenius, U., Tengzelius, U., & Ulfvengren, P. (2025). *Noise exposure from drone traffic: Scenarios and perspectives for analysis and control*. KTH Royal Institute of Technology.
- Rohrmeier, K., Wei, W., & Ison, D. (2025). Decoding the Vertiport: Planning for Urban Air Mobility. *Journal of Planning Literature*, 08854122251314481.
- Schweiger, K., & Preis, L. (2022). Urban air mobility: Systematic review of scientific publications and regulations for vertiport design and operations. *Drones*, 6(7), 179.
- SESAR. (2023). *U-space ConOps and architecture (edition 4)*.
- Shao, Q., Shao, M., & Lu, Y. (2021). Terminal area control rules and eVTOL adaptive scheduling model for multi-vertiport system in urban air Mobility. *Transportation Research Part C: Emerging Technologies*, 132, 103385.
- Yunus, F., Casalino, D., Avallone, F., & Ragni, D. (2023). Efficient prediction of urban air mobility noise in a vertiport environment. *Aerospace Science and Technology*, 139, 108410.

Biographical Notes

He Gao is a doctorate student at Chair of Land Management, Technical University of Munich. He previously received a PhD in Economics (2024) from Renmin University of China.

Prof. Walter Timo de Vries is the Chair and Professor of Chair of Land Management, Technical University of Munich.

Contacts

He Gao

Technical University of Munich
Arcisstrasse 21
80333 Munich
Germany
Email: he.gao@tum.de