TEISOND

Civic Infrastructure for Continuous Legitimacy Monitoring

WHITE PAPER

Version 3.0 (Public)

March 2026

Operated by Advanced Global Polling Technology Ltd (AGPT Ltd), United Kingdom

MISSION STATEMENT

Teisond provides citizens with permanent infrastructure to judge any official exercising governmental authority – making public legitimacy visible, continuous, and comprehensive – for every official, every day.

DOCUMENT STATUS

This White Paper presents the conceptual framework, methodology, governance principles, and implementation approach for Teisond – a civic technology platform for continuous legitimacy monitoring of government officials, operated by AGPT Ltd (UK)

This document is published for general informational purposes. It is intended for citizens, media organisations, researchers, civic organisations, and public officials interested in the Platform's mission and methodology For partnership inquiries: hello@teisond.com Website: teisond.com Language Versions

This White Paper is published in English as the canonical version. Official translations into national languages of EU member states are available on the respective national platform pages ({country}.teisond.com/whitepaper). In the event of any discrepancy between a translation and the English original, the English version shall prevail. Translation dates and update status are indicated on each language version.

COPYRIGHT AND USAGE

This document may be shared for informational and judgment purposes. Reproduction, distribution, or derivative works require written permission from AGPT Ltd This document is published in its entirety for public access

DISCLAIMER

This White Paper contains forward-looking statements regarding anticipated platform development, deployment timelines, and operational capabilities. Actual outcomes may differ due to technical, regulatory, and other factors The information provided is for informational purposes only and does not constitute legal advice or offer to sell securities Readers with specific legal or compliance questions are encouraged to seek qualified professional advice in their jurisdiction.

READER'S GUIDE

Document Scope

This is the complete public edition of the Teisond White Paper. It covers all ten sections plus four appendices. Some implementation details (anti-manipulation algorithm specifications, pricing schedules, vendor-specific configurations) are omitted by design to prevent gaming, not because they are restricted.

Document Structure

Section 1 – Introduction. What problem we solve, why existing instruments fall short, what Teisond is and is not, and why now.

Section 2 – Concept and Methodology. Core concept (legitimacy as a continuous signal), objects of monitoring, judgment mechanism, legitimacy index calculation, citizen value proposition, anti-manipulation safeguards, methodological transparency, limitations and biases.

Section 3 – Technical Architecture. Design philosophy, multi-tenant architecture, identity verification, data separation, publication controls, infrastructure resilience.

Section 4 – Revenue Model and Economics. Three-tier data access architecture, subscription model, primary and secondary revenue sources, cost structure, permanent exclusions.

Section 5 – Legal Structure and Jurisdictional Framework. Multi-entity architecture, GDPR as baseline, publication policy, contracting and dispute resolution, jurisdictional flexibility, contingency mechanisms, verification independence.

Section 6 – Governance and Ethics. Governing principles, roles and jurisdictions passport, ethical commitments and boundaries, stakeholder relationships, ethical dilemmas and resolution frameworks, governance evolution.

Section 7 – Roadmap and Implementation. Simultaneous EU presence strategy, wave launch plan, phased functionality, election sensitivity protocols, long-term vision, infrastructure and partner ecosystem, security and compliance roadmap, funding strategy and capital allocation.

Section 8 – Team and Organisation. Organisational philosophy, founder and leadership, team development, advisory board, governance and succession, community as institutional foundation.

Section 9 – Civic Ownership Architecture. The principle of citizen ownership, distributed infrastructure, token model and ownership pathway, governance distribution, transition from community formation to citizen ownership.

Section 10 – Conclusion. The case for continuous legitimacy monitoring, what the Platform achieves and does not, business model as mission alignment, governance as ongoing practice.

Appendices A–E. Theoretical and philosophical foundations; glossary; frequently asked questions; bibliography and references; governmental authority position estimates by level.

Navigation by Reader Type

Public officials: Executive Summary, Sections 1, 2 (overview), 4, 10. Time: 1-2 hours. Focus: what the Platform measures, how indices work, subscription services.

Civic organisations and partners: Executive Summary, Sections 1, 2, 5, 6, 7, 9, 10. Time: 3-5 hours. Focus: methodology, legal framework, governance principles, civic ownership architecture, deployment roadmap, partnership opportunities.

Academics and researchers: Sections 1, 2, 6, 9, 10, Appendix A. Time: 4-5 hours. Focus: theoretical grounding, methodology rigour, legitimacy concepts, civic ownership model, research opportunities.

Media and journalists: Executive Summary, Sections 1, 2 (overview), 6, 9, 10. Time: 1-2 hours. Focus: democratic innovation angle, citizen ownership model, governance principles, credibility indicators.

Investors and strategic partners: Executive Summary, Sections 4, 5, 7, 8, 9, 10. Time: 3-4 hours. Focus: revenue model, legal structure, roadmap, team, civic ownership trajectory, risk-return profile.

General public: Executive Summary, Section 1, Sections 9–10. Time: 30-60 minutes. Focus: what problem this solves, how it works, who owns it, why it matters.

Reading Time Estimates

Quick orientation: 30 minutes (Executive Summary only). Basic understanding: 1-2 hours (Executive Summary, Sections 1 and 10). Comprehensive grasp: 3-5 hours (all ten sections). Full reading: 6-8 hours (complete document including appendices). Academic study: 8-12 hours (with detailed attention to Appendix A).

Feedback Welcome

This White Paper will evolve through engagement with readers. If you find sections unclear, redundant, or insufficient, please provide feedback: hello@teisond.com

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The Accountability Deficit in Modern Democracy

In modern democracies, citizens elect representatives every four to five years but lack effective mechanisms to monitor them between elections. This temporal asymmetry – continuous authority, episodic oversight – is the Extra-Electoral Voter Influence Deficit.

Yet this represents only part of a larger accountability deficit. Beyond elected politicians, citizens encounter governmental officials daily – school principals, social services directors, police chiefs, tax office directors. These officials wield direct authority over citizens' lives, making decisions citizens must obey, yet no structured mechanism exists for citizens to hold them accountable. This is the Legitimate Public Influence Loop Deficit.

The scale of unmonitored authority is striking: for every official subject to sustained public scrutiny (500–2,500 nationally), there are 50 to 250 officials exercising comparable direct impact while operating in practical invisibility – totalling 50,000 to 500,000 officials exercising governmental authority in a typical democracy.

Together, these deficits reflect a single structural problem: the absence of mutual accountability infrastructure between citizens and officials. Officials exercise continuous authority; citizens lack any continuous, structured means to express judgment about how that authority is exercised. Existing mechanisms – elections, polls, petitions, protests, internal complaints, social media – each serve functions within their domains, but none creates continuous, comprehensive, reciprocal accountability (see Section 1 for detailed analysis).

Why Existing Mechanisms Fall Short

This is not a failure of effort but of structure. Elections bundle thousands of judgments into a single episodic choice. Opinion polls sample hundreds of respondents about a few hundred high-profile figures, leaving tens of thousands of officials unmeasured. Petitions and protests are episodic, high-cost, and low-signal. Social media generates noise without verification, structure, or privacy. Internal accountability mechanisms – ombudsmen, ethics commissions, inspector general offices – operate within institutional walls, invisible to the public they serve.

No existing instrument provides what democracy structurally requires: a permanent, verified, privacy-preserving channel through which any citizen can express judgment about any official, continuously, with results aggregated into transparent public indices. Teisond builds that missing instrument.

The Teisond Solution

Teisond is a civic technology platform addressing this deficit through a centralised multi-tenant system for continuous legitimacy monitoring. The Platform provides verified citizens with a permanent mechanism to judge any official exercising governmental authority in their country.

The defining innovation is universal scope. The Platform monitors not merely high-profile politicians but every individual holding power to make binding decisions in governmental capacity – from presidents and ministers to school principals, police chiefs, and social services supervisors. In a

mid-sized country (40–50 million population), the database covers approximately 40,000 officials across all four levels: all 640 national officials (100%), all 1,400 regional officials (100%), 18,000 municipal officials (mayors plus key councillors), and 20,000 local officials (20% coverage, focusing on high-interaction positions).

How It Works

For citizens: registration requires identity verification through commercial identity providers (document check + biometric liveness), ensuring one citizen equals one account. As national eID systems become available, the Platform integrates them as an additional verification path – but never depends on governmental infrastructure for its operation. Citizens express trust or distrust through a simple binary judgment – no justifications required. They can change or withdraw judgments as circumstances evolve. The process takes seconds.

For officials: the Platform calculates legitimacy indices – public, continuously updated percentages reflecting trust-to-total ratios. Officials monitor their own indices through subscription services (the Platform's primary revenue source), tracking real-time feedback and identifying areas requiring attention. A Right to Respond mechanism allows officials to publish statements linked to their indices – ensuring accountability runs in both directions.

For society: aggregated legitimacy data becomes mutual accountability infrastructure – accessible to media, researchers, civil society, and citizens. The data is structured, verified, continuous, and transparent – qualitatively different from episodic polls or unverified social media sentiment.

Privacy and Ethics First

The architecture embeds privacy by construction – not as policy that can be overridden, but as structural impossibility of misuse. Individual judgments never appear publicly – only aggregated indices. The system technically prevents political profiling: judgment histories are not stored, API endpoints return only aggregated statistics with k-anonymity thresholds, and the database schema excludes fields enabling demographic correlation of individual opinions. The question is not "will we protect privacy?" but "could anyone – including the Platform's own operators – violate it?" The architectural answer is no.

Creating a New Market Segment: Public Legitimacy Analytics (PLA)

The Platform operates within the Public Opinion Research & Social Insights market and expands it by creating a new segment – Public Legitimacy Analytics (PLA): continuous, citizen-sourced, aggregates-only measures of officials' legitimacy, published by office+period. Unlike episodic polling sampling a few hundred high-profile figures, PLA covers the full universe of offices and publishes continuously at near-zero marginal cost per additional user. PLA is a Blue Ocean move: it creates new demand rather than competing for polling share. The segment's products include the National Officials Legitimacy Index (NOLI), Office-Period Legitimacy Scorecards (OPLS), and Legitimacy Pulse & Trajectory with Risk Flags.

Citizen Value Proposition: Meeting Fundamental Psychological Needs

Platform viability rests not on civic duty appeals but on meeting fundamental human needs. Teisond addresses esteem needs through structured political participation that provides citizens with voice, tangible impact, and civic status unavailable through traditional mechanisms. When officials treat citizens dismissively, the Platform provides immediate recourse: a recorded judgment affecting the official's public legitimacy index – restoring dignity and agency where helplessness previously prevailed. This consumer-product approach creates sustainable engagement where abstract democratic appeals generate only temporary enthusiasm (see § 2.5 for comprehensive analysis).

Business Model and Revenue

Teisond operates as a two-sided platform. Citizens provide judgments at no cost. Officials, media, researchers, and consultants subscribe to access processed legitimacy data and analytical tools.

The primary revenue source is officials monitoring themselves – a psychologically universal motivation across all governmental levels. Whether presidents or school principals, officials care about their reputations. The motivations are identical across every level: real-time feedback on how their authority is perceived, unfiltered citizen sentiment unavailable through any other channel, peer pressure as colleagues begin monitoring their own indices, and professional necessity as legitimacy metrics become part of the landscape in which careers are built. Revenue is structurally aligned with mission: the very act of monitoring public legitimacy generates the data officials are willing to pay for. No advertising, no data sales, no grant dependence.

Legal Structure and Operational Architecture

Teisond is operated by AGPT Ltd, a UK-registered company (128 City Road, London EC1V 2NX). The UK jurisdiction provides sophisticated IP protections, the Defamation Act 2013's serious harm threshold for a platform publishing data about named officials, and globally recognised contract and corporate law.

AGPT Ltd operates a centralised multi-tenant architecture: a single codebase with country-specific configuration manages all EU deployments. AGPT Ltd acts as data controller in each jurisdiction, with citizen data stored locally within the EU – ensuring GDPR compliance independent of UK adequacy status. Each country's data is isolated; a breach in one deployment does not compromise others. Technical improvements benefit all countries simultaneously, while centralised operations eliminate network coordination overhead. If the Platform's growth justifies additional legal entities, the architecture accommodates them without requiring reconstruction – but the current structure is complete and operational as designed.

Founder and Team

Oleksiy Loboyko – Founder, CEO. Nearly five years developing the Teisond concept. Background in strategic communications, political analysis, and civic technology. Based in Ukraine; operational base transitioning to UK upon AGPT Ltd activation.

The founding stage is intentionally lean. The Platform's architecture – automated operations, config-driven country deployment, centralised multi-tenant design – is built to scale without proportional headcount growth. Core team recruitment (Technical Lead, Legal & Compliance Lead, Communications Lead) begins with launch-stage funding, prioritising mission alignment alongside technical capability. The organisational philosophy is automation-first: routine operations require no human intervention; people handle exceptions, strategy, and stakeholder relationships (see Section 8 for full team structure).

Implementation Timeline

All 27 EU countries receive national landing pages from day one – collecting registrations and signalling pan-European commitment. Full Platform activation proceeds in waves, determined by three criteria: verification infrastructure connected, AI-populated official database ready, and sufficient waitlist demand. Priority markets by infrastructure readiness and early demand include Estonia, Netherlands, Poland, Spain, and Germany – but the actual composition of each wave is determined by which countries meet activation criteria first, not by a predetermined schedule. The pace of expansion is constrained by operational readiness, not by engineering capacity.

Risks Acknowledged

Success requires sufficient citizen adoption, official subscription uptake across all levels, and resilience against manipulation attempts. Identity verification, anomaly detection, and privacy-by-construction architecture address technical risks. The centralised multi-tenant model ensures jurisdictional challenges in one country do not threaten the network. Alternative initiatives with better funding may fill the niche first. Mission drift under financial or political pressure remains a constant temptation. These risks are inherent in creating new civic infrastructure. Mitigation strategies reduce probability but do not eliminate uncertainty. The correct attitude is transparent acknowledgment, careful management, and continuous adaptation.

Long-term Vision

Teisond becomes standard democratic infrastructure – integrated into civic education and normalised as routine accountability mechanism. The Platform is designed from its first line of code for citizen ownership: an architectural commitment that the infrastructure ultimately belongs to those who generate its value, not to the entity that built it (Section 9). Mutual accountability between citizens and officials evolves from aspiration to operational reality. Citizens grow up expecting to judge officials continuously, and officials accept legitimacy monitoring as intrinsic to public service. The correct expectation: not revolutionary change but gradual, steady improvement in democratic accountability.

Call to Action

For Investors and Strategic Partners: This is civic infrastructure at the formation stage – a platform designed for the entire EU market with a clear revenue model, mission-aligned business logic, and a Blue Ocean positioning. The risk-return profile combines social impact with commercial viability. Read Sections 4, 5, § 7.8, and 8 to assess the opportunity.

For Officials: This is not a threat – it is a career management tool. Legitimacy indices give you what no other instrument provides: continuous, verified feedback from the citizens you serve. Early subscribers gain insight before public indices become widely cited. Read Section 2 and Appendix C (FAQ) to understand how the Platform works and what protections you have.

For Media: Move beyond the "ratings + scandals" paradigm. Replace assumptions with verified data. Tell political stories through the language of legitimacy. Make NOLI and office+period scorecards front-page metrics – a new system for public analytics.

For Researchers and Academia: Enter a new discipline at its formation stage. Make Public Legitimacy Analytics a living laboratory where theories of accountability and trust are tested on data. Set the academic standard for methodology in this field.

For NGOs and Civil Society: Engage as partner, user, and advocate. Support sustainable infrastructure of civic judgment – instead of investing in episodic bursts of media, petition, or street emotions. In the digital age, this is more effective, more reliable, and safer for participants.

For the Sceptical: Read on. This document is designed to withstand scrutiny, not to avoid it.

SECTION 1: INTRODUCTION

1.1 The Democratic Accountability Deficit

1.1.1 The Democratic Paradox and The "Incompetent Demos" Assumption

Modern representative democracy rests on a foundational bargain: citizens delegate power to officials in exchange for those officials' accountability to the public interest. In practice, however, much of democratic practice is built on a rarely stated premise: that most citizens are not capable of forming independent, consistent judgments about those in power and therefore must be managed through narratives, framing and targeted messaging. This implicit "incompetent demos" assumption is visible in how campaigns are designed, how opinion research is commissioned, and how institutions speak to the public.

Teisond starts from a different premise. It does not idealise citizens as perfectly informed or immune to manipulation, but it rejects the idea of an inherently incapable public. Instead, it treats civic capacity – the ability of ordinary people to observe what institutions are doing, compare this with their own expectations, and form a relatively stable opinion – as a widely distributed potential. Whether this potential is visible depends largely on the design of information systems and participation channels, not on the "quality" of the population.

When civic capacity is not given any stable, safe and recognisable outlet, democracies accumulate accountability deficits – manifesting as two interconnected problems.

1.1.2 The Extra-Electoral Voter Influence Deficit

Elections remain the primary formal mechanism through which citizens renew or withdraw mandates. Yet they occur infrequently – typically every four to five years. Between these episodic moments, elected officials continue exercising power while citizens lack any continuous, structured means to monitor or judge those actions.

This temporal asymmetry creates the Extra-Electoral Voter Influence Deficit: a structural disconnect between the continuous exercise of governmental power and the episodic nature of voter influence. Citizens may formally enjoy political equality yet experience practical powerlessness between elections. Traditional engagement mechanisms – town halls, public consultations, petition campaigns – show declining participation rates, as citizens perceive these tools as performative rather than effective.

1.1.3 The Legitimate Public Influence Loop Deficit

Beyond elected politicians lies a vast administrative apparatus of appointed officials exercising direct authority over citizens' day-to-day lives: A building inspector reviewing a permit application, with authority to approve, reject or delay a project. A school principal making disciplinary or curriculum decisions affecting a child's trajectory. A social services supervisor determining

eligibility and benefit levels for vulnerable populations. A police chief setting enforcement priorities and practices. A hospital administrator managing access to public healthcare services.

In each case, officials make decisions that citizens must obey. When these officials are incompetent, corrupt or arbitrary, citizens have very limited recourse. Internal complaint mechanisms prioritise institutional protection over citizen empowerment. Appeals are costly and time-consuming. Media rarely covers individual cases. Elections are irrelevant – these officials are never elected and never face voters directly.

This is the Legitimate Public Influence Loop Deficit: the near-total absence of citizen voice regarding the officials who most directly affect everyday life. A "legitimate public influence loop" means a structured, non-coercive feedback cycle that turns aggregated public judgments into visible signals and consequence-bearing incentives for officials.

1.1.4 The Scale of Unmonitored Authority

The numerical scale is striking. In a typical developed democracy:

Small democracies (3–5M population): ~150–300 officials routinely monitored; ~10,000–20,000 exercising governmental authority.

Medium democracies (10–30M): ~500–1,500 monitored; ~30,000–100,000 exercising authority.

A mid-sized EU member state (30–50M): ~1,500–2,500 typically monitored by traditional methods; ~100,000–150,000 exercising authority.

A large EU member state (50–100M): ~2,000–4,000 typically monitored; ~150,000–500,000 exercising authority.

For every official subject to sustained public scrutiny, there are 50 to 250 officials exercising comparable authority while operating in practical invisibility.

1.1.5 Two Deficits, One Problem

The Extra-Electoral Voter Influence Deficit and the Legitimate Public Influence Loop Deficit share a single root cause: the absence of mutual accountability infrastructure between citizens and officials. Officials exercise continuous authority; citizens lack continuous, structured means to express judgment about how that authority is exercised. The result is a structurally unidirectional relationship: authority flows downward continuously while the citizen's capacity to respond flows upward only episodically – and for the vast majority of officials, not at all. § 2.9 examines how continuous monitoring transforms this unidirectional relationship into mutual accountability.

1.1.6 Root Causes of This Deficit

Six structural factors explain why mutual accountability infrastructure has not emerged organically.

The most straightforward is technological. Until recently, no infrastructure existed to enable continuous, low-cost, verified, large-scale collection and publication of citizen judgments. Legacy tools – petitions, polling, elections – reflect the technological constraints of their eras and were

never designed to produce the kind of persistent, comprehensive signal that accountability infrastructure requires.

Institutional incentives work against external oversight. Governmental institutions resist mechanisms that create accountability pressure they do not control. Internal complaint systems are designed to protect institutional cohesion rather than empower citizens – and elected officials benefit from accountability gaps between elections, reducing any political incentive to build the infrastructure that would close them.

On the citizen side, collective action barriers prevent organic solutions. Individual citizens lack both the incentive and the capacity to organise comprehensive monitoring; where organised groups do emerge, they serve narrow interests rather than systematic oversight across all levels of governmental authority.

Two deeper factors compound these barriers. The "incompetent demos" assumption – the persistent presumption that citizens cannot form meaningful judgments about officials – discourages investment in participation infrastructure altogether. And the absence of a conceptual framework has meant that no standard concept of public legitimacy as a measurable, publishable construct has existed until now – leaving no demand signal for the infrastructure to measure it.

1.2 Existing Mechanisms and Their Limitations

Democratic societies have developed numerous accountability tools over centuries, each addressing specific aspects of the citizen-official relationship. Understanding why none of them – individually or collectively – fills the structural gap identified in § 1.1 is essential to understanding why Teisond exists. The limitation is not that these mechanisms are ineffective at what they do; it is that none of them does what is structurally missing: provide continuous, verified, universal, privacy-preserving citizen oversight of all officials.

1.2.1 Traditional Democratic Mechanisms

Elections remain the cornerstone of accountability, and rightly so – they are the only mechanism through which citizens can directly remove officials from office. But elections operate on multi-year cycles, produce binary outcomes (retain or replace) rather than continuous signals, and bundle thousands of policy and performance judgments into a single choice. A voter dissatisfied with their mayor's handling of housing but satisfied with their fiscal management cannot express this nuance through a ballot.

More fundamentally, elections apply only to elected officials – a fraction of those exercising governmental authority. The school principal, the social services supervisor, the police chief – none face voters. For the vast majority of officials who affect citizens' daily lives, elections provide zero accountability.

Referendums and recall mechanisms address specific situations but occur rarely, require high mobilisation thresholds, and apply only to elected officials in jurisdictions that permit them. Legislative oversight (parliamentary questions, committee hearings) operates within institutions rather than directly from citizens, concentrates on high-level policy, and rarely addresses individual official conduct. These are important mechanisms within their domains, but they do not create the continuous, universal feedback channel that is structurally absent.

1.2.2 Civic Participation Tools

Petitions signal citizen concern on specific issues but suffer from structural limitations as accountability mechanisms. They focus on policy demands rather than individual official performance. Signature thresholds are arbitrary – 10,000 signatures may trigger a response; 9,999 may not. Most critically, petitions are episodic: they arise in response to specific grievances and dissolve once the campaign ends, leaving no permanent infrastructure for ongoing oversight.

Protests effectively signal the intensity of public sentiment but are costly to organise, unsuitable for routine accountability, and necessarily concentrate on high-visibility targets. No one organises a street demonstration about a local tax office director's pattern of arbitrary decisions – yet that director's conduct may affect more citizens daily than a cabinet minister's policy choices.

Town halls and public consultations suffer from selection bias (attendees are unrepresentative), limited scope (one topic, one evening), performative dynamics (officials listen without obligation to respond), and geographic barriers. They create an appearance of participation without the infrastructure for sustained accountability.

1.2.3 Opinion Polling

Traditional polling has served democratic societies well as a snapshot tool, but it faces three structural limitations that make it inadequate as accountability infrastructure.

In coverage, polls focus on 500–1,500 high-profile figures at national level. The 50,000–500,000 local and regional officials who exercise direct authority over citizens' lives never appear in any poll – the Legitimate Public Influence Loop Deficit is entirely invisible to polling methodology. In cadence, polls are episodic snapshots: a monthly poll captures a moment but does not track the continuous evolution of public acceptance that legitimacy monitoring requires. And in verification, poll respondents are typically unverified – a single individual may respond to multiple polls, bots and coordinated campaigns can distort results, and the "one citizen, one account" principle that underpins democratic accountability has no enforcement mechanism in traditional polling.

Comprehensive, continuous, verified polling of all officials across all levels of government would cost tens of millions per wave – a practical and financial impossibility that ensures the accountability gap remains unfilled.

1.2.4 Digital and Social Media Tools

Social media platforms have transformed political expression, giving citizens unprecedented ability to voice opinions about officials. But as accountability infrastructure, social media suffers from fundamental design flaws.

There is no verification: anyone can create accounts, and coordinated inauthentic behaviour is endemic. There is no structure: a tweet criticising a mayor and a bot-generated pile-on are indistinguishable in the data. There is no aggregation: individual complaints do not compose into a meaningful measure of public acceptance. There is no privacy: citizens expressing political opinions on social media expose themselves to profiling, harassment, and retaliation.

Digital petition platforms (change.org, Avaaz) reduce friction compared to paper petitions but inherit the same structural limitations: episodic, issue-focused, unverified, and concentrated on high-visibility targets. Specialised civic tech platforms address specific needs – budgeting, local issue

reporting, legislative tracking – but none creates the comprehensive, continuous, verified judgment infrastructure that the accountability deficit requires.

1.2.5 Internal Accountability Mechanisms

Inspector general offices, ethics commissions, ombudsman institutions, and complaint review boards represent democracy's attempt to build accountability from within governmental structures. These institutions perform important functions – investigating misconduct, recommending reforms, mediating disputes – but they operate under structural constraints that limit their effectiveness as citizen-facing accountability infrastructure.

Institutional loyalty often takes precedence over citizen empowerment – complaint outcomes tend to favour the institution. Misconduct definitions are drawn narrowly, so that poor performance, arrogance, or systemic indifference may not qualify as actionable complaints. Processes are opaque: citizens file complaints into a black box and may never learn the outcome. And resource imbalances are structural – an individual citizen challenging an administrative decision faces the full weight of institutional legal resources, while sanctions, when they exist, are minimal and rarely public.

Most critically, internal mechanisms are invisible to the broader public. A citizen who files a complaint about a school principal contributes nothing to any public record of that principal's pattern of conduct. Each complaint is isolated; no aggregation occurs; no public signal emerges.

1.2.6 The Missing Mechanism

Across all domains, a specific combination of properties remains absent from any existing accountability tool. No current mechanism provides all of the following simultaneously:

It must operate continuously – permanent availability, not episodic campaigns. It must cover all officials exercising governmental authority, not only elected or high-profile figures. Participation must be verified – one citizen, one account, cryptographically enforced. Data must be structured – aggregated by office and time period into comparable indices – and publicly transparent, with results visible to citizens, officials, media, and researchers alike. Participation barriers must be low: seconds of engagement, not hours of organising. And privacy must be architectural: aggregates only, no individual profiling, no exposure of participants.

This is precisely the missing infrastructure for civic judgment that Teisond provides. The following section describes how.

1.3 The Teisond Solution

1.3.1 Core Concept

Teisond addresses the accountability deficit as a multi-tenant civic technology platform for continuous legitimacy monitoring, operated centrally by AGPT Ltd (UK). The Platform provides verified citizens with permanent infrastructure to judge any official exercising governmental authority in their country.

The defining innovation is universal scope – monitoring every individual holding power to make decisions binding on others in governmental capacity: "any person authorised to make decisions obligatory for others to implement." This includes national executives and legislators, regional and municipal authorities, appointed administrators, and local officials with direct citizen contact (building inspectors, school principals, police chiefs, social services supervisors, judicial officers, regulatory officials).

Across a typical mid-sized EU country, the Platform targets tens of thousands of officials across all four levels – a scale multiple orders of magnitude larger than traditional political monitoring, which typically covers 1,500–2,500 figures.

1.3.2 How It Works: Citizen Perspective

Registration requires verified identity authentication through commercial identity verification providers (document check + biometric liveness), ensuring one citizen equals one account. Where national eID systems are available and connected, they serve as an additional verification path. This verification is privacy-preserving: the Platform stores only a one-way cryptographic hash, never the citizen's name, ID number, or personal identifiers. Even Platform developers cannot determine who submitted a specific judgment.

Once verified, citizens search for officials by position, jurisdiction, or name. They view current indices, historical trends, and confidence intervals. The judgment itself is binary – trust or distrust – with no justifications required. Citizens can change or withdraw judgments as circumstances evolve, subject to rate limits preventing manipulation. The entire process takes seconds.

This low-friction design is deliberate. Civic participation tools that demand time, effort, or public exposure attract only the most motivated citizens. A mechanism that takes seconds and guarantees anonymity lowers the threshold to near-universal accessibility – including for citizens in sensitive positions who cannot afford to be seen criticising officials.

1.3.3 How It Works: Officials Perspective

The Platform calculates legitimacy indices – public, continuously updated percentages reflecting the ratio of trust judgments to total judgments, displayed with historical trends, confidence intervals, and comparative metrics. These indices are published for any office that meets minimum sample thresholds (typically 100 judgments); below that threshold, the Platform displays "Not enough judgments" rather than publishing potentially misleading numbers.

Officials interact with the Platform primarily through subscription services – the Platform's primary revenue source. Subscribers access detailed analytics: time-series data with hourly granularity, comparative benchmarks against peers at the same authority level, anomaly reports flagging unusual patterns, and early warning signals of legitimacy trajectory changes. A Right to Respond mechanism allows officials to publish statements linked to their indices, ensuring accountability is reciprocal.

Subscription pricing follows an accessibility-first approach with country-adjusted tiers by authority level (see Section 4). Critically, subscription status has zero effect on index calculation or publication – non-subscribing officials receive identical public indices as subscribers.

1.3.4 How It Works: Data Perspective

Privacy is architectural, not policy-based. The Platform records only current state (trust, distrust, or neutral) – never judgment histories. The database schema excludes fields enabling demographic correlation. API endpoints refuse individual-level queries regardless of authentication level. Political profiling is not merely prohibited by policy; it is structurally impossible because the data required for profiling does not exist in the system (see § 5.4 for comprehensive data governance).

Indices publish only when sample sizes exceed minimum thresholds (typically 100 judgments per office per period). Every published index includes a confidence interval reflecting sample size and variance. Below-threshold offices display "Not enough judgments" – protecting both statistical validity and participant privacy. Methodology is fully public and documented.

1.3.5 What Teisond IS and IS NOT

Teisond is a structured mechanism for continuous legitimacy monitoring across all governmental levels – infrastructure that supplements elections, media, and other accountability mechanisms rather than replacing any of them. It is a measurement system providing information that democratic society can incorporate into its existing processes: a permanent citizen voice that operates between and beyond elections; infrastructure designed for eventual citizen ownership, where the protocol belongs to those who use it (Section 9).

Equally important is what the Platform is not. It does not create binding legal obligations – officials with low indices retain full authority. It is not a political party or advocacy organisation – the Platform monitors all officials regardless of affiliation. It is not a social network or discussion forum – there is no commentary, no messaging, no content creation. And it is not a surveillance system – it monitors public acceptance of authority, not private behaviour; citizens are anonymous, officials are public figures exercising public power.

The distinction matters: Teisond does not tell officials what to do. It tells them – and the public – where they stand.

1.3.6 Theory of Change

Measurement creates accountability without legal force through several reinforcing channels. Officials who know their indices are public and persistent anticipate electoral consequences – low indices signal vulnerability long before voters reach the ballot box, creating incentive for responsiveness throughout the term. Beyond elections, indices become reputational capital: common knowledge shaping relationships with voters, parties, peers, and media. An official whose legitimacy index trends downward faces questions from colleagues, journalists, and constituents – even without any formal consequence.

Media amplification accelerates this dynamic. Journalists cite indices as standard reference points, creating feedback loops between coverage and Platform usage: a declining index becomes a news story; the news story drives more citizens to the Platform; more judgments refine the index further. Officials who subscribe to track their own indices generate the Platform's primary revenue while creating internal accountability pressure – the act of monitoring itself changes behaviour, because knowing that citizen sentiment is continuously visible makes unresponsive conduct harder to sustain. And competitive dynamics ensure that no official can simply ignore the data: rivals use declining indices as arguments for leadership changes, opposition parties cite them publicly, and within institutions legitimacy data informally shapes promotion, appointment, and assignment decisions – as reputation data always does, but now with a verified, public source.

1.3.7 Citizen Value Proposition

The Platform's viability rests not on civic duty appeals but on addressing a specific unmet psychological need: restoring dignity and agency in citizens' relationship with governmental authority.

When a municipal planning officer delays a permit without explanation, when a school principal ignores parental concerns, when a social services officer treats an applicant with contempt – citizens experience a specific combination of frustration, helplessness, and indignity. Traditional recourse is either unavailable (no complaint mechanism), ineffective (internal review protects the institution), or disproportionately costly (legal action, media campaigns, political organising).

Teisond provides immediate, low-cost, private recourse: a recorded judgment affecting the official's public legitimacy index. The citizen's experience is no longer invisible – it becomes part of a public signal. This restores agency without requiring confrontation, organisation, or public exposure. The psychological reward is concrete and immediate: "I was not powerless after all."

This consumer-product approach creates sustainable engagement where abstract democratic appeals generate only temporary enthusiasm. § 2.5 examines this value proposition comprehensively.

1.3.8 Public Legitimacy Analytics (PLA)

Teisond expands the Public Opinion Research market by creating a new segment – Public Legitimacy Analytics (PLA): continuous, citizen-sourced, aggregates-only measures of officials' legitimacy, published by office+period. Unlike polling that episodically samples ~500–1,500 high-profile figures, PLA covers 50,000–500,000 officials and publishes continuously at near-zero marginal cost per additional user.

PLA is a Blue Ocean move: it creates new demand rather than competing for existing polling share. Traditional polling answers "what do people think about the prime minister this month?" PLA answers "what is the current public acceptance of every official exercising governmental authority, from the president to the local school principal, updated continuously?" No existing product answers this question.

The segment's flagship products are the National Officials Legitimacy Index (NOLI), providing headline indices for all monitored officials; Office-Period Legitimacy Scorecards (OPLS), offering detailed cards per official per period with confidence intervals and trend data; and Legitimacy Pulse & Trajectory with Risk Flags, delivering time-series analytics with early warning indicators for significant shifts.

1.3.9 Limitations and Honest Expectations

Teisond does not claim to solve democratic accountability. It claims to provide one missing instrument – and that instrument has clear limitations.

Indices measure expressed confidence, not performance quality. Popular officials may implement destructive policies while competent technocrats face low scores. The index measures public acceptance of authority, not governance quality – and the WP makes this distinction explicit throughout.

Indices reflect participating users, not the entire citizenry. Participation skews are transparently disclosed: sample sizes, confidence intervals, and threshold notices accompany every published index. The Platform never claims representativeness it cannot demonstrate.

Indices capture direction but not intensity or reasoning. A citizen who profoundly distrusts an official and one who is mildly sceptical produce the same signal. This is a design choice – simplicity enables scale – but it means indices are blunt instruments, not precision diagnostics.

Indices create reputational pressure, not legal obligation. An official with 20% trust retains full legal authority. The Platform influences through visibility, not legal force – and some officials may simply ignore their indices.

The correct expectation is not revolutionary change but incremental improvement: making visible what was invisible, structuring what was chaotic, empowering what was powerless. Over time, this visibility changes norms – but the timeline is years, not months.

1.4 Why Now? Technological and Social Readiness

The accountability deficits described in § 1.1 are not new. The inability to build infrastructure addressing them is. Three convergent developments make Teisond viable now in a way that would have been impossible a decade ago.

1.4.1 Technological Enablers

Commercial identity verification infrastructure has reached the maturity, cost, and coverage necessary to guarantee one citizen = one account across all EU member states. Providers such as Veriff, Sumsub, and Onfido verify government-issued documents with biometric liveness checks in seconds, at scale, in every EU country – without dependence on governmental eID infrastructure. This is architecturally significant: civic accountability infrastructure must not depend on decisions by those it monitors. National eID systems (Estonia's ID-card, Poland's Profil Zaufany, Spain's Cl@ve, the Netherlands' DigiD) remain a welcome upgrade path where available – raising assurance levels and reducing per-user costs – but the Platform launches and operates independently of them. Without verified identity, any civic judgment platform degenerates into a manipulable poll. With it, the Platform can guarantee that every judgment represents a real, unique citizen – the democratic principle of 'one citizen, one account' enforced cryptographically.

Cloud infrastructure and API-first architectures have eliminated the capital barrier to serving millions of users across multiple jurisdictions simultaneously. A multi-tenant platform that would have required tens of millions in infrastructure investment fifteen years ago can now be deployed, scaled, and maintained at a fraction of the cost. Mobile-first progressive web applications remove installation friction – citizens participate through a browser, not an app store.

1.4.2 Social and Political Context

Declining trust in traditional institutions is not merely a polling finding – it is a lived reality shaping political behaviour across Europe. Citizens increasingly perceive formal participation channels (elections, consultations, petitions) as performative rather than consequential. This creates latent

demand for tools that provide genuine agency – not symbolic participation, but a mechanism whose outputs are visible, persistent, and publicly consequential.

Digital natives now constitute significant demographic cohorts across all EU member states. These citizens are comfortable with platform-based interaction, expect real-time feedback, and see no reason why civic participation should be confined to a polling booth once every four years. For this generation, the question is not "why would I use this?" but "why doesn't this already exist?"

Open data movements and transparency legislation have normalised the expectation that governmental performance should be measurable and public. Freedom of information regimes, open budget initiatives, and public procurement transparency have established the principle that citizens have a right to structured information about how authority is exercised. Teisond extends this principle from institutional performance to personal legitimacy – a natural next step that existing infrastructure has not yet taken.

1.4.3 Market Readiness

The multi-billion-dollar public opinion research industry faces structural disruption. Traditional polling – expensive, episodic, limited in scope, unverified in participation – is increasingly challenged by clients who demand continuous data, broader coverage, and methodological transparency. Teisond does not compete directly with Gallup or Eurobarometer; it creates an adjacent market segment (PLA) that serves needs polling structurally cannot address.

Political consultants, campaign strategists, and institutional analysts pay substantial sums for political intelligence. Legitimacy indices – continuous, verified, covering all governmental levels – provide a data product with no current equivalent. Media organisations seeking structured political data beyond episodic scandal coverage find natural integration points for legitimacy analytics.

Most importantly, the primary revenue source – officials subscribing to monitor themselves – taps a psychologically universal motivation that requires no market education. Officials at every level care about their standing with the public. The subscription does not need to be sold on abstract civic value; it sells itself as a career management tool. This alignment between mission and revenue is what makes the business model sustainable without grant dependence, advertising, or data monetisation.

SECTION 2: CONCEPT AND METHODOLOGY

2.1 Core Concept: Legitimacy as Continuous Variable

2.1.1 Rethinking Political Legitimacy

Political legitimacy – the recognition by citizens that those exercising governmental authority do so rightfully – has traditionally been treated as a binary attribute: a government either possesses legitimacy or it does not. Elections serve as the mechanism for conferring this status, and between elections, legitimacy is presumed until the next electoral test.

This binary treatment misrepresents the actual phenomenon. Legitimacy is not a fixed state conferred at a moment in time but a continuous, dynamic relationship between governed and governing. Every interaction between citizen and official either reinforces or erodes trust. A mayor who responds transparently to a local crisis strengthens legitimacy; one who evades accountability weakens it. These micro-interactions accumulate into collective sentiment that shifts continuously, yet traditional measurement captures only periodic snapshots.

Teisond reconceptualises legitimacy along three axes. First, legitimacy is subjective and personalised – it exists in the judgment of individual citizens, not as an objective institutional property. There is no legitimacy "by definition"; there exists only the aggregate of individual trust assessments. Second, legitimacy is continuously variable – it fluctuates in response to official actions, policy outcomes, personal experiences, and information exposure. Third, legitimacy is measurable – if defined as aggregated citizen trust toward specific officials, it can be captured, quantified, and published through appropriate infrastructure.

This reconceptualisation has practical consequences. If legitimacy is binary, measurement is unnecessary – elections suffice. If legitimacy is continuous, measurement becomes essential – societies need instruments tracking how legitimacy evolves continuously, across the full spectrum of officials exercising authority – including the vast majority who never face elections at all. The Platform provides this instrument.

2.1.2 Legitimacy as "Currency of Expectations"

Legitimacy functions as a currency of expectations between citizens and officials. Citizens "invest" expectations in officials through the act of delegating authority. Officials "spend" this currency through their decisions and conduct. Fair, competent governance maintains or increases the balance. Corruption, incompetence, or dismissiveness depletes it.

Unlike financial currency, legitimacy currency has no central bank controlling supply. Each citizen holds their own reserve of expectations and distributes trust independently. The Platform makes this distributed currency visible by aggregating individual assessments into public indices – creating, for the first time, a transparent "exchange rate" between citizen expectations and their acceptance of how authority is exercised.

This framing has practical implications for officials: legitimacy is not an inexhaustible resource. Officials who consistently deplete citizen trust face declining indices visible to voters, supervisors, media, and peers. The Platform does not create this dynamic – it makes it measurable and visible.

2.1.3 From Episodic Sampling to Continuous Census

Traditional political measurement relies on episodic sampling: polling organisations survey 1,000–2,000 respondents at intervals of weeks or months, extrapolating to populations of millions. This captures roughly 0.001–0.01% of the population per survey and covers only a few hundred high-profile figures.

Teisond inverts this model. Rather than sampling a fraction of citizens about a fraction of officials, the Platform enables the entire verified citizenry to judge any official at any time. If participation reaches even a fraction of the voting population, the resulting dataset dwarfs any polling sample – and covers tens of thousands of officials across all levels of authority rather than a few hundred national figures. Updates are continuous rather than periodic, identity verification ensures data integrity impossible in anonymous polling, and the marginal cost of additional participants and officials approaches zero once infrastructure exists.

This is not incremental improvement on polling methodology but a categorical shift: from episodic sampling to continuous census of expressed political judgment.

2.1.4 Early Warning System and Accountability Between Elections

Continuous legitimacy measurement creates an early warning system for democratic governance. Gradual erosion of an official's legitimacy index signals emerging problems before they escalate into crises. A mayor whose trust index declines steadily over several months receives a clear signal that citizen confidence is eroding – providing opportunity for course correction rather than discovering the problem only at the next election.

For appointed officials who never face elections, this early warning function is even more critical. A regional administrator whose legitimacy index shows persistent decline generates a signal visible to supervisors, media, and the public – creating accountability pressure where no structured mechanism existed before.

The Platform thus addresses a structural absence in democratic accountability: elected officials govern without structured feedback between elections, and appointed officials – the vast majority of those exercising governmental authority – operate without any such feedback at all.

2.1.5 Limitations and What Legitimacy Indices Don't Measure

Legitimacy indices measure expressed citizen trust, not objective governance quality. High legitimacy does not guarantee competent governance; low legitimacy does not prove incompetence. Popular officials may implement destructive policies while maintaining public support; competent officials making difficult decisions may face low indices.

Indices reflect participating citizens, not the entire population. If participation is unevenly distributed, indices may overrepresent certain demographics. The Platform mitigates this through transparency: publishing sample sizes and confidence intervals alongside every index, so that readers can assess the statistical reliability of the data they are interpreting.

Indices capture direction (trust/distrust) but not reasoning. A citizen distrusting an official due to corruption and one distrusting due to policy disagreement register identically. The Platform measures what citizens express, not why – leaving interpretation to the democratic ecosystem of media, researchers, and public discourse.

These limitations are inherent in any measurement system and are disclosed transparently rather than concealed. The Platform provides information; democratic institutions and citizens determine its significance.

2.2 Objects of Monitoring

2.2.1 Inclusion Principle and Definition

Teisond monitors individuals holding governmental authority – those making decisions that affect citizens' lives and claiming to act on citizens' behalf. The inclusion principle is universal: any person authorised to make decisions obligatory for others to implement in governmental capacity. This definition captures the essence of governmental authority: the power to impose binding decisions on others. Whether that authority derives from election, appointment, professional qualification, or administrative hierarchy, if an individual exercises it on behalf of government, they become an appropriate subject for citizen judgment. This principle is the Platform's core criterion, distinguishing it from all existing accountability mechanisms.

The definition excludes individuals who do not exercise governmental authority regardless of their public influence. Business leaders, activists, and public figures may shape public discourse but do not make decisions binding on others in a governmental capacity. Career civil servants who implement policies established by others – without discretionary decision-making authority of their own – fall outside the scope for the same reason.

Certain categories of officials who do hold governmental authority require additional safeguards to protect institutional independence. Judicial officers, for example, are included in the monitoring framework because they exercise binding authority over citizens – but the Platform draws a clear boundary between measuring citizen confidence in an office-holder and creating pressure on specific rulings. The safeguards applicable to judicial and other independence-sensitive positions are described in the Note on judicial monitoring (§2.2).

2.2.2 Four Levels of Governmental Authority Across the EU

Teisond operates across 27 EU member states with a combined population of approximately 450 million citizens. Each country has a unique institutional structure – federal or unitary, centralised or devolved, with varying numbers of administrative tiers – but the underlying pattern of governmental authority is universal. The Platform accommodates this diversity through country-specific configuration (§5.3.1), not structural redesign: the same methodology, privacy architecture, and publication standards apply to every political system.

The Platform identifies four levels of governmental authority, each containing officials whose decisions bind citizens. The levels reflect organisational structure and jurisdictional scope, not importance or hierarchy of accountability. A municipal planning officer exercises authority over citizens' lives just as surely as a cabinet minister – the scope differs (individual permits versus national policy) but the principle of accountability applies equally.

Across the EU-27, the total number of officials exercising governmental authority at all four levels is estimated at 2–4 million individuals. Traditional political monitoring typically covers only a few hundred high-profile elected officials per country. The Platform extends accountability infrastructure to officials operating at every level of governmental administration. The revenue implications of this scope – why officials at every level become natural subscribers – are described in Section 4.

2.2.3 Level 1: National Authority

National-level officials exercise authority affecting the entire country's population – broad policy, national legislation, executive action, and institutional direction. This level encompasses heads of state and government, cabinet ministers and junior ministers, members of national legislatures, senior political appointees heading major agencies, and – where appointments are political – supreme or constitutional court justices and heads of regulatory bodies.

A typical EU member state has 250–800 national-level officials; across the EU-27, the total is approximately 12,000–15,000. These are the highest-visibility political actors – the only tier covered by national polling, major media, and traditional political monitoring. The Platform supplements this existing ecosystem with continuous rather than episodic measurement. A detailed breakdown of Level 1 positions and estimated counts is provided in Appendix E.

2.2.4 Level 2: Regional Authority

Regional-level officials exercise authority over provinces, states, autonomous regions, or other first-tier administrative subdivisions. The scope of regional authority varies significantly across the EU: federal systems (Germany, Austria, Belgium) have strong regional governments with substantial legislative and executive powers; quasi-federal systems (Spain, Italy) grant significant devolved autonomy; unitary centralised systems (France, Poland) maintain regional administration that primarily implements national policy. This level includes regional heads of government, regional legislators, and senior administrators at the provincial or intermediate tier where applicable.

A typical EU member state has 400–2,500 regional officials, with federal systems at the higher end. Regional officials occupy a middle visibility tier – they receive coverage in regional media and face some polling before elections, but systematic continuous monitoring rarely extends comprehensively to this level. The Platform closes this gap. A detailed breakdown is provided in Appendix E.

2.2.5 Level 3: Municipal Authority

Municipal-level officials exercise authority over cities, towns, and local administrative units. This level shows the greatest structural variation across the EU – from Denmark's 98 consolidated municipalities to France's 35,000+ communes – but the principle is uniform: mayors, municipal councillors, and senior municipal administrators make binding decisions affecting citizens' daily lives.

The accountability deficit is sharpest at this level. Media coverage and political analysis typically reach perhaps 50–200 prominent mayors nationally. For the remaining thousands of municipal leaders – and tens of thousands of councillors – no continuous accountability mechanism exists between elections. They exercise governmental authority while operating outside any structured public scrutiny. A typical mid-sized EU country has 25,000–80,000 municipal officials; across the

EU-27, the total is an estimated 800,000–2,000,000. The Platform adapts to each country's municipal structure through configuration, not code changes. A detailed breakdown is provided in Appendix E.

2.2.6 Level 4: Local Officials and Appointed Administrators

This level represents the Platform's most significant innovation: creating accountability infrastructure for officials who previously operated with virtually none beyond internal institutional processes. Level 4 encompasses officials exercising direct authority over citizens in daily life – school principals, police chiefs, building inspectors, hospital directors, social services supervisors, tax office directors, court judges, and administrators across every sector where governmental decisions affect individual citizens more immediately than abstract policy debates.

The scale is substantial. A typical mid-sized EU country has 50,000–150,000 Level 4 officials; across the EU-27, this is by far the largest category – an estimated 1.5–3 million individuals. These officials are largely invisible to traditional accountability mechanisms: no polling organisation surveys citizen trust in a local building inspector, no media outlet tracks the legitimacy of a regional hospital director. Yet these are the officials whose decisions citizens experience most directly and most frequently.

Judicial officers are included at this level because they exercise governmental authority – but the Platform draws a clear boundary. Legitimacy indices for judicial officials measure citizen confidence in the office-holder, not agreement or disagreement with specific rulings. The Platform publishes no case-level data and provides no mechanism for citizens to respond to individual decisions. Judicial independence remains paramount; monitoring measures public acceptance without creating instruments that could influence outcomes.

Database population for Level 4 begins with the operational team identifying officials from public records and government rosters, prioritising the most common citizen touchpoints: schools, police, building inspections, social services, hospitals, and tax offices. AI agents progressively expand coverage as the Platform scales. A detailed breakdown by sector is provided in Appendix E.

2.2.7 Implementation Approach

The MVP scope is defined by completeness at a specific authority level, not by partial coverage across all levels. Teisond launches with a fully populated database of Level 1 officials (§2.2.3) across all active countries.This ensures that citizens in every country can begin rendering judgments on the officials with the highest public visibility from the first day of operation. The database is pre-populated by the Platform's operational team using official government registries, parliamentary rosters, and publicly available institutional records. Expansion to Levels 2–4 follows as country deployments mature and operational capacity allows – each level added only when the database for that level meets completeness and accuracy standards across the relevant jurisdiction.

Once launched, the database is maintained through continuous monitoring of official sources: electoral results, government reshuffles, term expirations, and organisational changes are tracked and reflected in the database as they occur. Expansion toward comprehensive Level 2–4 coverage – potentially encompassing hundreds of thousands of officials per country – is driven by AI agents that systematically identify and catalogue office-holders from public records, government rosters, and institutional directories. This automation is what makes pan-European coverage at all four authority levels operationally feasible without proportional growth in team size.

2.2.8 Exclusions and Boundary Cases

The exclusion of private citizens and career civil servants without decision-making authority is described in §2.2.1. Two further categories are clearly excluded: monarchs and ceremonial heads of state whose functions are constitutionally symbolic, and officials of international organisations who do not exercise authority within a national governmental framework.

Several boundary cases require case-by-case assessment. Semi-autonomous agencies are included where they exercise binding authority over citizens – a national energy regulator issuing mandatory compliance orders qualifies; a purely advisory economic council does not. Public university leaders are included because they administer publicly funded institutions with authority over students and staff; their counterparts at private universities are excluded. Professional associations are included only where they exercise delegated governmental authority with legal force – a bar association that controls admission to legal practice qualifies; an industry lobby group does not. Advisory bodies are excluded unless they hold enforcement authority. The treatment of judicial officials – included with specific independence safeguards – is addressed in §2.2.6.

The guiding principle across all boundary cases is consistent: does this person make decisions that citizens must obey? If yes, they are an appropriate subject for civic judgment. If they advise, recommend, or process without deciding, they are excluded.

2.3 Judgment Mechanism

2.3.1 The Binary Choice: Trust or Distrust

Citizens judge officials through a deliberately simple mechanism: trust or distrust, with neutrality expressed by not judging at all. The point is to keep the cognitive cost close to zero so anyone can participate within seconds, including right after an encounter with an official, without composing justifications or learning a rating rubric. This simplicity reflects how people naturally form judgments about authority in real life.

A binary input avoids false precision. Ten-point or multi-criteria scales look scientific but invite arbitrary distinctions (what really separates a "6" from a "7"?) and force hidden weighting schemes. A clear yes/no signal, aggregated at the office+period level, produces transparent percentage indices that the public, media, and officials can read consistently, with confidence intervals rather than spurious decimals.

Privacy improves when the Platform asks less. A single trust/distrust signal reveals minimal information about the person who submitted it; multi-dimensional ratings or written reasons would create richer profiles and greater risk of political-opinion exposure. Keeping inputs binary fits the aggregates-only design and helps meet data-protection expectations without sacrificing usefulness.

The binary choice does not flatten nuance. Citizens can trust some officials and distrust others; they can change or withdraw their current judgments as circumstances evolve. At the collective level, nuance emerges through movement over time and across offices: steady declines, sharp drops with recovery, or volatile trajectories all carry different meanings that a complex input scale does not magically improve.

In short: a small, human-scale action from each person becomes a clear, auditable public signal when aggregated – without over-engineering the input or over-exposing the individual.

2.3.2 What Is NOT Collected

The binary judgment mechanism is deliberately minimal not only in what it asks but in what the Platform retains. No reasons or justifications are collected. No intensity of feeling is measured. No policy-specific breakdowns, forced rankings, or demographic correlations are captured. No per-user judgment histories or behavioural profiles are derived. The Platform stores only the current state – citizen X currently trusts or distrusts official Y – and nothing more.

This minimisation is architectural, not merely a policy commitment. The database schema omits fields for the categories listed above; API endpoints cannot expose what does not exist; queries below publication thresholds are rejected. Even in a hypothetical data breach, there is no trove of personal political opinions to extract – because the Platform never collected them. The canonical list of excluded data categories and the enforcement mechanisms that make this irreversible are described in §5.4.7; publication thresholds in §5.4.8; and API access controls in §5.4.10.

2.3.3 User Actions and Constraints

The Platform offers three actions to every verified citizen. A citizen can submit an initial judgment – trust or distrust – for any official in the database. A citizen can change that judgment as their assessment evolves, moving from trust to distrust or vice versa. And a citizen can withdraw a judgment entirely, returning to neutral and removing their input from the aggregate calculation. Participation is voluntary and self-directed: a citizen may judge many officials or only a few, and may register without judging anyone at all.

Three constraints protect the integrity of the signal. Each citizen holds one judgment per official at a time – trust, distrust, or neutral – with no mechanism for multiple simultaneous positions. Judgments cannot be delegated to a party, influencer, or organisation; each judgment is an individual act, preventing organised influence-trading or mass delegation. And no bulk operations exist – there is no mechanism to apply a single judgment across a group of officials, ensuring that each input reflects a deliberate, individual assessment.

Rate-limiting mechanisms constrain how frequently a citizen can change a judgment for any given official, preventing rapid toggling that would add noise rather than signal. The specific parameters of these controls are not published, consistent with the anti-manipulation disclosure policy described in §5.4.11.

2.3.4 The Role of Neutrality

Every citizen begins in a neutral state toward every official. Neutrality means only that no judgment has been submitted – it carries no positive or negative signal. A citizen may be neutral because they lack information about an official, have had no recent interaction with that office, or have deliberately chosen to abstain. The Platform does not interpret silence. Indices are calculated from the ratio of active judgments only: if an official has received 10,000 trust and 5,000 distrust judgments, the Trust Index is 66.7% – regardless of how many registered citizens have not judged that official. Non-participants are excluded from the denominator entirely.

This design reflects a methodological principle: treating silence as a signal would inject noise into every index. No citizen can meaningfully track thousands of officials, and requiring blanket judgments would produce arbitrary inputs rather than genuine assessments. The Platform measures explicit, voluntary acts of civic judgment – trust or distrust – and nothing else. Neutrality is both a right and a quality safeguard: it preserves citizen autonomy, protects data integrity, and

ensures that every published index reflects intentional expression rather than assumptions about what silence might mean.

2.4 Legitimacy Index Calculation

2.4.1 Basic Formula

For each official, the Platform maintains two live counters: the number of verified citizens who currently express trust (TI) and the number who currently express distrust (DI). These are integers that update as citizens add, change, or withdraw judgments. No per-user histories are stored. The published Legitimacy Ratio (LR) is the share of trust among all active judgments, scaled to 0–100: LR = TI / (TI + DI) × 100. The Distrust Ratio is the simple complement: 100 – LR. Both figures are displayed together to aid interpretation. Citizens who have not submitted a judgment – or who have withdrawn one – do not enter either counter and do not affect the denominator; the ratio reflects the balance of expressed trust versus distrust among those who chose to judge (§2.3.4).

LR is published only when the total number of judgments exceeds the country's minimum publication threshold. Below that threshold, the Platform displays "Not enough judgments" rather than a potentially unreliable figure. Above it, LR is rounded to avoid false precision and reduce re-identification risk, and is always shown alongside its confidence interval (§2.4.2). The full publication policy – including k-anonymity requirements and threshold enforcement – is described in §5.4.8–5.4.9.

LR is a privacy-preserving, office-level signal of current public acceptance of authority. It is not a measure of job performance, policy wisdom, or legal validity, and it is not an election forecast. It should always be read together with its confidence interval, sample size, and any applicable transparency notes.

2.4.2 Confidence Intervals and Sample Size

Every published Legitimacy Ratio is displayed with a 95% confidence interval. The interval reflects statistical uncertainty and depends on both the observed ratio and the sample size n (where n = TI + DI). Larger samples narrow the interval; smaller samples widen it.

For calculation purposes, the trust proportion is expressed as p = TI/n (a decimal between 0 and 1), then scaled to the 0–100 display format. For mid-range proportions and sufficiently large samples, the 95% confidence interval uses the normal approximation: $CI_{95} = p \pm 1.96 \sqrt{[p(1-p)/n]}$. When n is small or p lies near 0 or 1, the Platform uses Wilson intervals, which remain well-calibrated at boundaries where the normal approximation breaks down. To illustrate: an official with 1,550 trust and 950 distrust judgments (n = 2,500) would show LR = 62.0 with a 95% CI of approximately 60.1–63.9. An official with only 120 total judgments and a near-even split would show a much wider interval – roughly ±9 percentage points – signalling that the index has not yet stabilised.

The Legitimacy Ratio is displayed at one decimal place and the confidence interval bounds at one decimal place, avoiding false precision. Results are published only when n exceeds the country-specific minimum threshold; below it, the Platform displays "Not enough judgments" rather than an unreliable figure. Overlapping confidence intervals between two officials signal that apparent differences may be statistical noise rather than meaningful divergence – a point the Platform makes visible through its presentation design.

2.4.3 Temporal Aggregation and Trend Analysis

Legitimacy indices are not snapshots but time series. The Platform computes indices over configurable windows – a 7-day short horizon and a 30-day default horizon – and preserves historical values for longitudinal analysis. Values are recalculated on a regular cadence, and the current index is displayed distinctly from the historical curve so that the latest value is clear without losing trajectory context. Publication thresholds apply to every plotted point: only above-threshold aggregates appear; sub-threshold windows display "Not enough judgments." Each historical point carries uncertainty consistent with the confidence interval for its window.

When a material recalculation or methodology correction occurs, the historical record is versioned and accompanied by a dated public note explaining what changed and why. No historical value is silently revised. This versioning discipline ensures that any reader – journalist, researcher, or official – can fully reconstruct the history of any index and verify that the data they cited at an earlier date has not been retroactively altered.

2.4.4 Comparative Metrics

A single official's index gains meaning in context. The Platform supports four comparative lenses: peer comparison (officials in the same type of role – all cabinet ministers, all mayors in a region), geographic comparison (officials across regions within a country, surfacing territorial patterns), temporal comparison (an official's own trajectory over time), and institutional comparison (aggregated averages for departments, agencies, or authority levels, revealing systemic strengths or weaknesses). Cross-jurisdiction benchmarking – comparing equivalent roles across countries – is possible but requires careful interpretation given differences in political culture and participation rates. In every case, comparisons display sample sizes and confidence intervals, and sub-threshold entities are excluded until they meet publication minimums.

2.4.5 Handling Edge Cases and Anomalies

The Platform favours disclosure over suppression. When anomalies are detected – coordinated campaigns, sudden volume spikes, or statistically improbable patterns – the index continues to be published but is accompanied by a visible flag and a plain-language transparency note. The detection framework and governance principles behind this approach are described in §5.4.11. No index is silently removed or adjusted; readers always see both the data and the context needed to interpret it.

The Platform also handles lifecycle transitions for the officials it monitors. Newly added officials may show volatile indices while sample size builds; results are published only once the threshold is met, and early above-threshold indices are marked as preliminary. When an official leaves office – through resignation, term expiry, or electoral loss – the profile is preserved as a historical record with no further judgments accepted, maintaining the accountability trail without distorting active data. Where a single individual holds distinct positions carrying separate authority, each position maintains its own index. Changes in party affiliation or personal name do not reset an index; accountability follows the incumbent's tenure in the office, not the party label.

2.5 Citizen Value Proposition: The Psychology of Participation

2.5.1 The Participation Paradox

Democratic theory assumes engaged citizens continuously monitoring governmental authority. Reality diverges sharply. In established democracies, most citizens demonstrate minimal political engagement outside electoral campaigns – they cannot name their representatives, rarely follow proceedings, and invest virtually no time tracking official performance.

Traditional civic technology initiatives have repeatedly attempted to bridge this gap – parliamentary monitoring platforms, petition systems, civic complaint mechanisms – yet struggle to achieve sustained mass engagement. The pattern repeats: initial enthusiasm, modest early adoption, gradual decline into niche usage by politically hyperactive minorities. The failure stems from fundamental misunderstanding of user motivation: most civic technology treats political participation as moral obligation, appealing to civic duty and democratic responsibility. These appeals prove insufficient against platforms engineered to deliver immediate psychological rewards.

Anthony Downs's theory of "rational ignorance" explains why: the personal cost of becoming informed exceeds the marginal benefit of one voice among millions. Traditional civic platforms fail because they demand cognitive investment without offering commensurate individual return. The social benefit may be substantial, but individual incentive remains weak.

The participation paradox thus defined: democracy requires citizen engagement, yet citizens rationally decline to engage because individual costs exceed individual benefits. Any platform attempting to bridge the accountability vacuum must solve this paradox through genuine individual value creation, not stronger moral appeals.

2.5.2 The Need: Dignity, Agency, Voice

Every person who lives under governmental authority experiences a fundamental tension: others make decisions that bind them, yet they possess no routine mechanism to signal whether they accept that authority as legitimate. This is not a gap between elections – it is a permanent condition. Elections, where they occur, offer a single binary choice once every several years, covering only a fraction of the officials who exercise power. Appointed administrators – the officials citizens encounter most frequently and most directly – never face even this minimal test. The result is a chronic asymmetry: authority flows downward continuously, while the citizen's capacity to respond flows upward only in rare, episodic, and highly constrained moments.

This asymmetry frustrates one of the most basic human needs. In Abraham Maslow's hierarchy, esteem needs occupy the level above safety and belonging – and in the EU's social democracies, where institutional infrastructure broadly secures those lower tiers, esteem needs become an active motivator for large segments of the population. Esteem operates on two axes. Internal esteem – self-respect, agency, the sense that one's judgment matters – is threatened whenever a person feels powerless or insignificant in the face of authority. External esteem – recognition, social standing, the knowledge that others see and hear you – is threatened when no channel exists for expressing that judgment visibly. Governmental authority, exercised daily without structured feedback, frustrates both. The citizen is not merely uninformed or disengaged; the citizen has no instrument through which engagement could take meaningful form.

2.5.3 The Product: From Need to Market

This is the need the Platform addresses – not as a civic obligation but as a consumer product that delivers individual psychological value. Participation on Teisond is a small, low-cost action – a single trust or distrust judgment – that produces an immediate, visible result: a published index that moves. The citizen's voice enters a public signal that officials, media, and peers can see. Internal esteem is restored through agency – the knowledge that one's judgment has been registered and counted. External esteem is restored through visibility – the knowledge that the aggregated expression, to which one has contributed, is now a matter of public record. No existing instrument offers this combination. Petitions collect signatures but produce no persistent signal. Opinion polls sample a fraction of the population and disappear after publication. Social media posts generate noise without structured measurement. The Platform converts individual judgment into a continuous, measurable, and publicly visible index – a mechanism that did not previously exist. Participation is sustained by intrinsic motivation, not artificial incentives: the Platform deliberately avoids gamification (points, badges, rewards) because people who find an activity inherently meaningful continue it longer than people who are externally incentivised to perform it. Participation aligns with natural attention cycles – a political scandal, a bureaucratic encounter, a media report – rather than demanding dedicated time or artificial engagement rituals.

The consumer who uses this product is any adult resident of an EU member state who has ever felt that their opinion of an official does not matter. This is not a niche audience of political activists; it is a general population whose esteem needs are routinely unmet by existing democratic infrastructure. The market is the 450 million residents of the EU-27. The realistic addressable segment – citizens motivated enough to complete identity verification and submit at least one judgment – is estimated at 20–40% of the voting population in active countries, a penetration rate that vastly exceeds any existing civic technology platform.

Teisond is, in this sense, a social enterprise operating through market mechanisms. The citizen's need for voice and recognition generates participation. Participation generates data. Data generates revenue from officials who subscribe to monitor their own indices (Section 4). Revenue sustains the infrastructure that serves the citizen's need. This cycle is self-reinforcing: the more citizens participate, the richer the data, the more valuable the product, the more sustainable the mission. The business model does not merely align with the social mission – it is structurally identical to it.

2.5.4 From Individual Motivation to Collective Accountability

Teisond's democratic value does not reside in any single judgment. It emerges from the aggregation of thousands of individual acts into a collective signal that neither any participant nor the Platform itself deliberately constructs. Each citizen judges an official for personal reasons – frustration with a bureaucratic encounter, recognition of competent leadership, a need for voice. No individual judgment carries public weight on its own. But as judgments accumulate past statistical thresholds – minimum sample sizes, confidence intervals, publication criteria – the index becomes a continuously updated public signal. The transition is not gradual but threshold-based: invisible until statistically robust, then visible to everyone simultaneously. This pattern is familiar from other domains: markets emerge from individual transactions, encyclopedias from individual edits, election outcomes from individual preferences. The Platform applies the same aggregation logic to the relationship between citizens and governmental authority – a space where no equivalent mechanism currently exists.

The relationship between individual motivation and collective effect is reinforcing. As participation grows, indices become more statistically robust, which increases their public visibility, which increases officials' responsiveness, which demonstrates to citizens that participation matters, which motivates further participation. The individual act is identical at every stage of this cycle; its collective significance is transformed. What distinguishes this aggregation from existing mechanisms is the combination of four properties simultaneously: it is continuous (not episodic),

verified (one citizen, one account), census-aspiring (not sampled), and citizen-initiated (not researcher-driven). The Platform's role is not to create accountability but to create the conditions under which accountability emerges from citizen behaviour that is individually motivated and collectively significant. The structural dynamics this produces – reciprocity, mutual accountability, systemic pressure – are examined in §2.9.

2.5.5 A New Channel Between Governing and Governed

The Platform does not appeal to idealistic motivations. It does not rely on civic consciousness, democratic duty, or altruism. Its reason for existence is simpler: to open a channel of information exchange between groups whose relationship has been structurally unidirectional. Governmental authority flows downward continuously – decisions, permits, refusals, regulations, enforcement. The reverse signal – how citizens perceive and judge those who exercise this authority – has no channel. It is not suppressed; it is architecturally absent. Elections offer a momentary, highly constrained valve; for appointed officials, no valve exists at all. The result is a latently conflictual relationship in which frustration accumulates without outlet, and officials operate without structured feedback from the people over whom they hold power.

Citizens participate because they want to be citizens, not subjects – to have a voice in a relationship where voice has been structurally denied. Officials engage because they are the interested recipients of what the Platform produces: a continuous, public signal of how their authority is perceived. Media process the Platform's data because they cannot afford to remain outside the social context it creates. Each participant in this information exchange acts from their own motivation. The cumulative effect is societal acceptance of a civic technology infrastructure that none of them individually set out to create or initiate.

Personal interests and societal interests need not conflict. The Platform, as a new social institution, generates value that exceeds any individual participant's intent: mutual recognition between governed and governing, a continuous legitimacy signal where none existed, and reduced tension in relationships that were previously one-directional and opaque. This added social capital – accountability, transparency, bilateral awareness – is not the Platform's direct function but the emergent consequence of an information channel that connects two groups whose exchange was, until now, blocked.

2.6 Anti-Manipulation Safeguards

Even with verified identity ensuring one citizen, one account, organised groups can coordinate legitimate users to express trust or distrust en masse. This is not fraud – all participants are real citizens with verified accounts – but organised mobilisation, a form of democratic participation no different in principle from petition drives or voter turnout campaigns. The methodological challenge is that unchecked coordination could make indices primarily reflect organisational capacity rather than organic citizen sentiment. In the most serious scenario, foreign states could activate networks of authentic but directed users, producing movements in indices that carry no domestic democratic meaning. The Platform's response is transparency rather than censorship: coordinated activity is detected, flagged, and disclosed to readers – never silently suppressed or removed. The detection architecture and disclosure framework are described in §3.5 and §5.4.11.

2.7 Methodological Transparency and Auditability

2.7.1 Open Methodology

The Platform's methodology for calculating legitimacy indices is fully public. The formulas, confidence interval methods, aggregation rules, publication thresholds, edge case handling, and exclusion criteria are documented and available for scrutiny by researchers, journalists, regulators, and citizens. What is not published – deliberately – is the specifics of anomaly detection algorithms, because disclosing detection parameters would provide a roadmap for circumventing them. The methodology documentation is versioned: when any element changes, the previous version is preserved and a dated public note explains what changed and why (§2.4.3). The Platform's public API (§5.4.10) enables independent verification of published indices by researchers, journalists, and any interested party.

2.7.2 Independent Audits and Third-Party Verification

The Platform's transparency commitment extends to independent verification of its own operations. Security, methodology, privacy, and compliance are subject to audit by third parties – with scope and frequency staged to the Platform's maturity. The roadmap progresses from pre-launch code review and vulnerability assessment through independent penetration testing to formal certification as operational scale requires (§7.7). Audit results are published, with sensitive security details redacted to prevent exploitation.

2.8 Limitations, Biases and Ethical Considerations

2.8.1 Self-Selection Bias

Participation is voluntary. Citizens who register and submit judgments are not a random sample of the population – they are likely to be more civically engaged, more digitally connected, and may differ from the general population by age, education, or location. Legitimacy indices therefore reflect the judgments of participating citizens, not a census-level portrait of public sentiment. This is an inherent property of the Platform, not a flaw to be engineered away.

The Platform addresses self-selection through three measures. First, transparency: every published index is accompanied by its sample size and confidence interval, and the Platform's documentation makes explicit that indices represent aggregated signals from participants, not weighted estimates of population-level opinion. Second, sustained effort to broaden participation through multilingual interfaces, accessible design, and outreach to under-represented communities as operational capacity allows. Third, honest refusal to apply demographic weighting or representativeness corrections – the Platform does not collect demographic attributes by design (§5.4.7), and applying statistical adjustments to data it does not possess would be methodologically dishonest.

2.8.2 Digital Divide and Access Barriers

Not all citizens have equal access to digital platforms. Older adults, rural residents, low-income communities, and people with disabilities face real barriers to participation, and the Platform's indices will reflect these gaps until they are addressed. Two commitments are non-negotiable from launch: citizen participation is permanently free, and the Platform is designed for accessibility by default – mobile-first delivery, multilingual interfaces covering all official languages in each country, and compliance with established accessibility standards. Broader inclusion measures – partnerships with public institutions, low-bandwidth optimisation, digital literacy support – are developed progressively as operational capacity and local partnerships allow. The Platform does not claim to have solved the digital divide; it commits to not deepening it.

2.8.3 Populism and "Tyranny of the Majority" Concerns

Legitimacy indices can reflect popular prejudices, short-term emotion, or majoritarian bias against officials who protect minorities. The Platform acknowledges this tension. It measures what citizens express; it does not declare who "deserves" trust. Low indices never remove an official from office –constitutional processes remain intact. Trends over time and peer comparisons help readers interpret whether movement reflects bias or performance.

Does making real-time sentiment visible incentivise pandering? Does continuous monitoring invite a "tyranny of the majority," where officials shy away from defending unpopular minorities? The pressures toward populism already exist through polls, media, and social platforms; Teisond organizes those pressures and makes them legible. Indices are not commands: officials may act against the grain when they judge it right, and sustained explanation can rebuild trust over time. Democracy is majoritarian but constitutionally tempered; Teisond does not alter that balance. Minority protections come from constitutions, courts, and rights frameworks, not from any index. The Platform reports a single dimension – popular legitimacy – and leaves other dimensions visible to the institutions designed to uphold them.

Over the long run, transparency creates counter-pressure to demagoguery. Persistent pandering without results tends to erode trust; difficult decisions that pay off can restore it, and the trajectory is visible in the record.

In short, Teisond makes public sentiment continuous and interpretable without turning it into a veto. It clarifies – rather than replaces – the existing constitutional settlement in which rights, courts, and law protect minorities while elected and appointed officials remain free to lead, explain, and, when justified, swim against the current.

2.8.4 Manipulation by Misinformation

Citizens will sometimes judge officials on the basis of false claims or partisan propaganda. The Platform neither can nor should audit the reasoning behind each judgment. Democracy does not require voters to demonstrate adequate information at the ballot box, and the Platform extends that principle between elections. What the Platform measures is perception: whether people currently accept an official's authority. That is a feature, not a flaw, because legitimacy concerns belief, not a forensic score of "deservingness."

Safeguards focus on integrity and clarity rather than censorship. The anomaly detection and disclosure framework (§2.6, §5.4.11) ensures that unusual activity is visible to readers. Over time, misinformation-driven swings usually wash out, while durable sentiment persists; trajectories make that distinction visible. The Platform remains neutral: indices report aggregated citizen sentiment,

not claims about truth or merit. Interpreters may discount a move if they judge it misinformation-driven, but the Platform itself does not adjudicate narratives.

Combating falsehoods is the job of journalism, fact-checking, civic education, and open debate. The Platform's role is to make sentiment visible, auditable, and researchable – without exposing individuals or enabling profiling. Sharp moves should be read as prompts to seek context; the Platform supplies the visibility, the public sphere supplies the argument.

2.8.5 Emotional vs. Rational Judgment

Citizens often judge officials in the moment – after a specific encounter, a news story, or a visible decision. The Platform does not try to engineer this away. Democracy already runs on a mix of reason and affect; continuous measurement simply makes that mix visible. In political behaviour research, emotion is not a flaw to be filtered out – it is part of how people assess trustworthiness, fairness, and competence. When citizens feel angry, disappointed, reassured, or inspired, those are politically meaningful signals.

Continuity moderates volatility. Because measurement is ongoing, short-lived surges of emotion are followed by recovery if concern fades; sustained anger shows up as a lasting decline. The same mechanism that records a sharp drop also records correction when it happens. Officials are not powerless in the face of emotional response – clear communication, remedial action, or policy adjustment can rebuild confidence over time. Continuous feedback creates opportunities for course correction rather than locking anyone into a permanent verdict.

The alternative is not pure rationality. Elections themselves are saturated with emotion – campaigns, rallies, partisan identity, last-minute scandals. If continuous legitimacy monitoring is criticised for reflecting emotional responses, consistency requires the same criticism of elections. The Platform's contribution is to turn democratic sentiment into a transparent time-series signal that can be read with context, confidence intervals, and historical trajectory, rather than guessed at from episodic impressions.

2.9 From Unidirectional Authority to Mutual Accountability

The methodology described in preceding sections delivers more than measurement infrastructure. It creates a structural transformation in the relationship between citizens and governmental authority. This final section addresses the question that binds the methodology together: what does continuous legitimacy monitoring actually change?

2.9.1 The Structural Problem

The relationship between governing and governed is structurally unidirectional. Authority flows downward continuously – decisions, regulations, enforcement, resource allocation – while the citizen's capacity to respond flows upward only through episodic, constrained mechanisms. Elected officials face voters once every four to five years; between elections, the feedback channel is effectively closed. For appointed officials – the vast majority of those exercising governmental authority – even this episodic mechanism is absent. The asymmetry is not incidental; it is architectural. Elections, courts, media, and internal oversight each address fragments of the accountability relationship, but none creates what is fundamentally missing: a continuous channel

through which citizens' aggregated judgments flow back to the officials whose authority depends, in democratic theory, on public acceptance. This blocked feedback defines the relationship as latently conflictual – frustration accumulates without outlet, officials operate without structured signals from those over whom they hold power, and both sides default to assumptions about the other rather than information. The Platform opens this channel.

2.9.2 How Continuous Monitoring Creates Reciprocity

For the citizen, the opening of this channel transforms a fundamental experience. Where previously frustration, approval, or concern had no structured outlet, each now has a destination: a judgment that enters a public, continuously updated index. The citizen does not gain power over the official's decisions – constitutional authority remains intact – but gains something that did not exist before: a persistent, visible signal that flows upward. The relationship shifts from unidirectional to bilateral. The citizen is no longer a silent recipient of authority but a participant in an information exchange whose output is public.

For the official, the channel introduces structured feedback from a direction that was previously silent. Officials routinely receive feedback from superiors, institutional peers, and media – but almost never from the citizens over whom they exercise daily authority. The Platform changes this. An official who subscribes to monitor their own index develops continuous awareness of how their authority is perceived by those it affects. Declining indices prompt investigation; stable indices provide confirmation. This is not external pressure imposed from above – it is information arriving from below, through a channel that previously did not exist.

For the system as a whole, the cumulative effect is a shift in norms. When legitimacy indices become available across officials at every authority level, comparison becomes inevitable – political rivals reference them, media highlight contrasts, supervisors notice patterns, and citizens develop expectations of responsiveness. The pressure this creates is not coercive; no index removes anyone from office. It is informational: the cost of ignoring citizen sentiment rises steadily as the signal becomes more visible, more cited, and more expected. Over time, responsiveness to the governed becomes not an act of virtue but a practical response to an environment in which indifference is publicly measurable.

2.9.3 Mutual Accountability as Emergent Property

What emerges from this bilateral information flow is not a single outcome but a cluster of social capital that did not previously exist. Accountability is the most visible component – officials become responsive to citizen sentiment because that sentiment is now public and persistent. But it is not the only one. Mutual recognition arises: officials see, for the first time, how their authority is perceived by those it affects, while citizens see that their judgment registers and matters. Reduced tension follows: relationships that were latently conflictual because feedback was blocked become less volatile when both sides have access to a shared signal. None of these outcomes is the Platform's direct function. The Platform opens an information channel between governing and governed; the social capital is generated by the exchange that flows through it. This is not revolutionary transformation but incremental structural change – adding a missing feedback mechanism to democratic systems that currently operate without one. The theoretical significance of this transformation – including principal-agent theory, information asymmetry resolution, and compound republic theory – is examined in Appendix A.

2.9.4 Integration with Existing Democratic Mechanisms

Continuous legitimacy monitoring supplements existing democratic mechanisms – it does not replace or compete with them. Elections remain the source of legal mandate. Courts protect constitutional rights. Media investigate and inform. Civil society advocates for change. Each addresses a distinct dimension of the democratic relationship. What the Platform adds is a specific capability absent from all of them: continuous, structured, verified expression of citizen judgment across every level of governmental authority, flowing through a channel that remains open between elections and extends to officials who never face elections at all. As the Platform's data becomes established, it naturally integrates into the existing ecosystem – informing media coverage, enriching public discourse, and providing researchers with a longitudinal signal that no current instrument produces.

2.9.5 From Methodology to Significance

The methodology described in this section defines what the Platform measures and why. The sections that follow describe how this methodology is implemented in practice: the technical architecture that enforces its privacy and integrity guarantees (Section 3), the revenue model that sustains it without compromising neutrality (Section 4), the legal framework that governs its operation across 27 jurisdictions (Section 5), and the governance principles that ensure it remains faithful to its mission as it scales (Section 6).

SECTION 3: TECHNICAL ARCHITECTURE

This Section provides a non-technical overview of the Platform's architectural principles. It is intended to demonstrate that Teisond's design enforces the privacy, neutrality, and transparency commitments described in Sections 2 and 5 – not merely as policy, but as structural properties of the system itself. Detailed technical specifications are maintained internally; what follows describes the architectural decisions that shape the Platform's guarantees to citizens, officials, and regulators.

3.1 Design Philosophy: Single Engine, Multiple Configurations

Teisond operates as a multi-tenant platform: a single codebase serves all 27 EU country deployments simultaneously. Each country runs on its own isolated configuration and database, but shares the same core logic for index calculation, anti-manipulation safeguards, and data governance.

This architecture means that a methodological improvement or security patch deployed once is instantly active across all countries. There is no version fragmentation – every national deployment runs identical code at all times. National differences – language, governmental structure, authority levels, identity verification method, publication thresholds – are expressed entirely through configuration files, not through code changes. Launching a new country deployment requires no software modification: it requires a configuration file describing that country's institutional structure, verification provider, data residency location, and localisation parameters (§5.3.1).

The multi-tenant approach also eliminates a class of governance risks inherent in federated architectures – version fragmentation, methodological divergence between deployments, and inconsistent access controls. Because AGPT Ltd operates all deployments centrally from a single codebase, privacy protections, methodology, and incident response are uniform across every jurisdiction by construction, not by agreement.

3.2 Identity Verification: One Citizen, One Account

The integrity of legitimacy indices depends on a strict one-citizen-one-account guarantee. Without it, the Platform would be indistinguishable from an online poll – vulnerable to bots, duplicate accounts, and coordinated manipulation. Teisond achieves this guarantee through identity verification against government-issued documents with biometric liveness confirmation – provided by commercial identity verification infrastructure (Veriff, Sumsub, or equivalent providers operating across all EU member states). No weak verification methods (email, phone, social media) are accepted. Where national eID systems (eIDAS framework) are available and connected to the Platform, they serve as an additional verification path – raising assurance levels and reducing per-verification costs. However, the Platform's operational independence requires that its launch and

continued operation never depend on governmental identity infrastructure. Civic accountability infrastructure must not be architecturally dependent on decisions by the officials it monitors.

Upon verification, the citizen's identity is immediately transformed through irreversible cryptographic hashing (SHA-256 with a seasonal salt). The Platform retains only the resulting fingerprint – never the original identity data. The salt rotates periodically, rendering previous fingerprints unlinkable to current identities. After a defined retention period, old salts are permanently deleted, making retroactive identification technically impossible – not merely prohibited by policy, but precluded by the absence of the cryptographic material required to reverse the hash.

The result: the Platform can confirm that each judgment comes from a unique, verified citizen – without knowing or being able to discover who that citizen is. This is a stronger privacy guarantee than pseudonymisation (where re-identification remains theoretically possible) and a stronger integrity guarantee than any unverified system can provide.

3.3 Data Architecture: Separation and Minimisation

The Platform enforces architectural separation between three categories of data: Identity data (hashed fingerprints confirming uniqueness) is stored in isolation from judgment data (trust/distrust records). These datasets cannot be joined or cross-referenced, even by system administrators. Technical controls – separate storage systems, no shared keys, no join endpoints – enforce this separation at the infrastructure level, not through access policies that an administrator could override.

Judgment data records only current state: trust, distrust, or neutral. No history of changes is stored. When a citizen updates their judgment, the previous value is overwritten, not appended. There is no timeline of opinions, no pattern of changes, no behavioural data that could be used for profiling.

Aggregate index data (published legitimacy ratios, trends, confidence intervals) is computed from judgment data but contains no individual-level information. The aggregation layer applies k-anonymity thresholds before any data reaches publication endpoints, ensuring that no published aggregate permits inference of an individual judgment.

The database schema is deliberately minimal. Fields enabling demographic correlation, political profiling, or behavioural tracking do not exist in the schema – they cannot be collected because the containers for such data were never built. This is what "privacy by construction" means in practice: not a policy restraining what the system could do, but an architecture limiting what the system can do.

3.4 Publication and Access Controls

Legitimacy indices are published only when sample sizes exceed configurable minimum thresholds, set per country and authority level. Below these thresholds, no data – not even qualitative indicators – is displayed. This protects both statistical validity (small samples produce unreliable indices) and participant privacy (small groups may be re-identifiable). Confidence intervals accompany every published index, and indices are rounded to prevent false precision.

These are not presentation choices – they are privacy mechanisms embedded in the publication pipeline. The full publication policy, including k-anonymity requirements and anomaly flagging, is described in §5.4.8–5.4.9.

Access to published data follows a three-tier architecture – anonymous qualitative indicators, registered citizen indices, and subscriber analytical data – described in §4.1. Access controls are enforced at the infrastructure level: no API endpoint exists that could return individual-level data regardless of authentication, and queries requesting demographic breakdowns or sub-threshold disaggregation are rejected by the system architecture, not by a policy layer that could be circumvented (§5.4.10).

3.5 Anti-Manipulation Safeguards

The Platform's identity verification requirement (§3.2) eliminates the most common attack vector against online platforms: bot accounts and duplicate registrations. Since only document-verified citizens with biometric liveness confirmation can participate, the threat model shifts from fabricated identities to coordinated campaigns by real people – a fundamentally different and more constrained problem. Commercial identity verification providers add further layers: cross-session face matching to prevent one person registering with multiple documents, device fingerprinting during verification, and known-fraud database checks.

For coordinated campaigns by verified citizens, the Platform employs multiple detection layers – including volumetric, temporal, and geographic analysis – combined with rate-limiting mechanisms that constrain the speed at which any campaign can move an index. Detailed specifications are not published to prevent gaming. When anomalies are detected, the Platform favours disclosure over suppression: affected indices display a visible anomaly indicator with a plain-language note explaining what was detected. The index continues to be published, but readers can see that the data may reflect organised activity rather than organic citizen sentiment. The governance framework for this approach – including the principle that legitimate collective action should not be silently erased – is described in §5.4.11.

3.6 Infrastructure Resilience

The Platform is designed so that no adverse event in a single jurisdiction can disable operations elsewhere. Country-level tenant isolation ensures that a legal challenge, regulatory action, or technical failure affecting one country's deployment does not propagate to others. Each country's data resides in jurisdiction-appropriate infrastructure within the EU, hosted by providers selected for GDPR compliance, availability guarantees, and independence from the governmental structures the Platform monitors.

Contingency infrastructure matures alongside the Platform itself. At launch, resilience rests on tenant isolation, automated backups, and the ability to migrate any country's data to an alternative hosting provider within the same jurisdiction or the EU. As operational scale grows, the contingency architecture expands to include provider diversification, failover testing, and documented migration procedures. The staging of these measures follows the security and compliance roadmap described in §7.7; the governance framework for contingency decisions – including trigger conditions and escalation procedures – is described in §5.7.

The resilience philosophy is simple: no single point of failure – technical, legal, commercial, or political – should be capable of disabling the Platform or compromising its data.

SECTION 4: REVENUE MODEL AND ECONOMICS

This Section explains how Teisond sustains itself financially while preserving its public mission. The core principle: the Platform's revenue model must never create incentives that conflict with accurate, neutral measurement of legitimacy. Every design choice described below is tested against this principle – if a revenue source could bias indices, distort methodology, or compromise privacy, it is excluded regardless of its financial attractiveness.

4.1 Data Access Architecture: Three Tiers

Teisond provides three levels of access to legitimacy data, each serving a distinct purpose and a distinct audience:

Tier 0 (Anonymous, no registration): Visitors see a qualitative indicator – a categorical description of the overall legitimacy position – without numerical values.This tier satisfies basic public curiosity, demonstrates the Platform's existence and scope, and serves as a conversion funnel: visitors who see the qualitative indicator and want to know the actual number are prompted to register. Tier 0 costs the Platform nothing to serve and creates organic demand for Tier 1.

Tier 1 (Registered citizen, free): After identity verification, citizens see the Legitimacy Ratio (LR) as a percentage, a trend indicator, and the current sample size. This level is sufficient for meaningful civic oversight and media citation. It supports the Platform's public mission – every citizen can see where every official stands – and remains permanently free. Tier 1 is not a loss leader; it is the Platform's core product for citizens, and its existence is what generates the judgment data that makes the commercial product valuable.

Tier 2 (Subscriber, paid): Subscribers access the full analytical layer: absolute trust and distrust counts, consensus metrics, time-series data, comparative analytics across officials at the same authority level, and anomaly detection reports.This is the Platform's commercial product – and the only tier that generates direct revenue.

This three-tier structure ensures that the public interest in transparency is served unconditionally (Tiers 0 and 1), while the commercial value of deep analytics funds the Platform's operation (Tier 2). The architecture makes the mission and the business model mutually reinforcing: the more citizens participate (free), the richer the data becomes, the more valuable the subscription product grows. The publication policy governing what data is available at each tier, and what is never published, is described in §5.4.9.

4.2 Primary Revenue: Officials Monitoring Themselves

The Platform's primary revenue source is subscription fees from officials who monitor their own legitimacy indices. This is not a speculative revenue hypothesis – it rests on a psychologically universal mechanism observable at every level of governmental authority.

The psychological mechanism is straightforward: once an official's legitimacy index becomes publicly visible and is cited by media, colleagues, and constituents, ignoring the Platform becomes costlier than subscribing. Media organisations and opposition figures may have access to detailed analytics; an official without a subscription sees less about their own standing than observers do. This information asymmetry creates natural demand for the subscription product. The motivations are identical whether the official is a prime minister or a municipal council member: control (real-time feedback), curiosity (unfiltered citizen sentiment), peer pressure (colleagues monitoring creates social expectation), and professional necessity (indices become metrics that matter in career advancement and institutional standing).

Critically, a subscription has zero effect on the official's index. The subscription purchases access to a mirror – it does not alter the reflection. This separation is absolute and architecturally enforced. A subscribing official and a non-subscribing official with identical citizen judgments will display identical public indices.

Subscription pricing is fixed by authority level (not by judgment volume or index value), eliminating any potential conflict of interest between Platform revenue and judgment activity. Pricing is country-adjusted for accessibility – an L4 local official in Estonia and an L4 local official in Germany pay amounts calibrated to local economic conditions, but both pay a fixed fee that does not vary with their index. Pricing scales with authority level – from L4 at the lowest price point to L1 at the highest.

The pricing structure is designed so that no official is priced out of self-monitoring. If legitimacy indices become a standard reference point in democratic discourse – cited by media, consulted by voters, referenced by institutional peers – then the subscription is not a luxury but a professional necessity. The Platform's growth strategy depends on making this transition: from novelty to standard infrastructure.

4.3 Secondary Revenue Sources

The Platform generates secondary revenue from three subscriber categories, all accessing Tier 2 data (§4.1) through the same analytical layer but for different professional purposes.

The most natural secondary audience is news media. For newsrooms under constant cost pressure, the Platform replaces ad-hoc polling commissions with a continuous, structured data source. A journalist with subscriber access can generate a data-driven story on any official at any time – time-series trends, comparative analytics across officials at the same authority level, anomaly indicators – without commissioning a poll or waiting for an election cycle. The value proposition is not a single dataset but a permanent editorial resource.

Academic institutions and policy organisations represent a second subscriber category, accessing anonymised and aggregated datasets for democratic governance research under terms that prohibit re-identification attempts. Early research partnerships serve a dual purpose: they generate

peer-reviewed validation of the Platform's methodology and contribute to the emergence of Public Legitimacy Analytics (PLA) as a distinct research domain.

The third category comprises political consulting firms, public affairs agencies, and governmental analytics units. These subscribers use legitimacy indices as inputs to their own advisory work – campaign strategy, stakeholder mapping, institutional risk assessment. As the Platform's data becomes a standard reference in the political intelligence ecosystem, revenue from this segment grows proportionally. The governance framework for institutional and consultancy relationships is described in §6.4.6.

4.4 Cost Structure: Automation-First Operations

Teisond operates on a lean, automation-first model. The Platform is designed so that routine operations – index calculation, publication, anti-manipulation monitoring, country configuration updates, threshold enforcement – run autonomously without human intervention. Human resources are reserved for exceptions, strategic decisions, quality oversight, and stakeholder relationships.

This architectural choice has direct financial implications. Adding a new country deployment does not require proportional headcount growth – it requires deploying a configuration file and activating identity verification for that jurisdiction. The marginal cost of an additional country is primarily infrastructure (hosting, verification integration), not personnel. This means the Platform's cost structure scales sublinearly with geographic expansion, creating improving unit economics as the Platform expands to additional countries.

The automation-first approach also reduces operational risk. Routine decisions – when to publish, what thresholds to apply, how to flag anomalies – are made by the system according to documented rules, not by individual operators exercising discretion. This consistency is both an operational advantage and a governance commitment: the Platform operates the same way whether it is monitoring a controversial politician or an obscure local administrator.

4.5 What the Revenue Model Excludes

The following revenue sources are permanently excluded by design – not as aspirational commitments subject to future reconsideration, but as structural constraints embedded in the Platform's architecture and governance framework.

The first category of exclusions eliminates revenue streams that would compromise the Platform's neutrality. Advertising of any kind – display ads, sponsored content, paid visibility – is excluded because it creates incentives to maximise engagement rather than accuracy. Sale of user data or behavioural profiles is excluded because individual-level data does not exist in exportable form by architectural design, and even if it did, selling it would destroy the trust that makes participation possible. Political consulting and campaign services are excluded because offering strategic advice to officials would compromise the neutrality that makes the data credible. The Platform provides data; it does not advise on how to use it.

The second category eliminates any arrangement where payment could create differentiated treatment. A subscribing official's public profile appears identically to a non-subscriber's –

subscription purchases analytical access, not favourable presentation. Pricing is fixed by authority level and jurisdiction; no official, institution, or third party can pay to alter, suppress, delay, or prioritise the publication of any index. There is no premium tier that buys better visibility and no volume discount that rewards institutional pressure.

These exclusions answer a question every informed reader will ask: "If this platform depends on officials' money, how do I know the indices are honest?" The answer is architectural: the revenue model is designed so that paying more – or paying at all – cannot change what the Platform publishes. The mirror does not flatter its subscribers.

SECTION 5: LEGAL STRUCTURE AND JURISDICTIONAL FRAMEWORK

5.1 Operational Structure: Centralised AGPT Ltd

Teisond operates through a single legal entity. AGPT Ltd, incorporated in the United Kingdom, is the sole operator, sole data controller, and sole intellectual property holder for the Platform across all 27 EU member states. There are no franchised national operators, no local subsidiaries, and no delegated controllers.

This single-entity structure is a deliberate choice for the current stage of the project. It eliminates coordination overhead between entities, removes transfer pricing complexity, ensures that methodology and privacy standards are applied uniformly without inter-entity negotiation, and places undivided accountability for every country deployment in one organisation. When something goes wrong in any jurisdiction, the responsible party is unambiguous: it is AGPT Ltd.

The legitimate concern with a single-entity model is concentrated risk – legal challenge, regulatory action, or operational failure in one jurisdiction could in principle affect the operator's capacity to serve others. This risk is mitigated through three mechanisms: country-first data residency ensures that each country's data is architecturally isolated regardless of corporate structure (§5.3); insurance coverage addresses defamation, cyber, and professional liability exposure across jurisdictions (§5.5.3); and contingency mechanisms define the conditions and procedures for maintaining service continuity under adverse scenarios (§5.7). If the Platform's growth reaches a stage where a second legal entity – such as an EU-domiciled subsidiary – becomes operationally necessary, the path for structural evolution is described in §5.9.

5.2 Jurisdictional Home

5.2.1 Jurisdictional Choice: Why the United Kingdom

AGPT Ltd is incorporated in the United Kingdom. This is a deliberate jurisdictional choice grounded in the specific risk profile of the Platform, not a default selection based on founder convenience.

The most consequential factor is UK defamation law. The Defamation Act 2013 introduced a serious harm threshold and other reforms that rebalanced the relationship between reputation protection and public interest speech. For a platform that publishes legitimacy data about named officials across 27 countries, the legal environment governing its home jurisdiction matters more than for most technology companies. English law provides a more predictable and balanced framework for this type of publication than many EU civil law jurisdictions, where reputation protections can be significantly stronger.

English contract and corporate law offer a second practical advantage. The legal frameworks governing AGPT Ltd's Terms of Service, data processing agreements, and commercial contracts are well-understood internationally, reducing friction in relationships with EU-based partners, subscribers, and regulatory counterparts. UK intellectual property law provides clear and enforceable protections for the Platform's software, methodology, trademarks, and trade secrets.

The United Kingdom also benefits from an EU adequacy decision under GDPR Article 45, which provides the legal foundation for AGPT Ltd to act as data controller for EU citizen data (§5.4.3). This adequacy status, combined with Standard Contractual Clauses maintained as an independent safeguard, ensures that the Platform's jurisdictional home does not create a data protection deficit for EU participants.

5.2.2 Legal Structure

AGPT Ltd is incorporated as a private company limited by shares under the UK Companies Act 2006. This corporate form provides limited liability protection for the founder and any future shareholders, straightforward governance through standard articles of association, and full transparency through Companies House filing obligations. The structure accommodates future evolution – including the admission of mission-aligned investors or transition toward community governance models – without requiring re-incorporation.

The path for such evolution is described in §5.9. AGPT Ltd also holds all intellectual property associated with the Platform, including the Teisond trademark, software, methodology, and visual identity.

5.3 National Data Governance

5.3.1 Country-Specific Configuration

Each country deployment is operated by AGPT Ltd through a country-specific configuration file that defines all jurisdiction-specific parameters without requiring separate legal entities or local teams. The configuration determines four core dimensions of each deployment.

Identity verification method is defined per country. Each country configuration specifies which verification provider is used for citizen registration. The default method is a commercial identity verification provider (such as Veriff or Sumsub) offering document scanning with biometric liveness checks – a method that works uniformly across all 27 member states without dependency on national infrastructure. Where a national eID system is available and connected to the Platform – such as Estonia's X-Road, Poland's Profil Zaufany, or Spain's Cl@ve – it is offered as an additional verification option that raises assurance level while preserving the Platform's operational independence from any single government-controlled system.

Official database scope determines which categories and levels of officials are monitored in that country, sourced from official registries and institutional directories. The scope expands progressively – from Level 1 national officials at launch to comprehensive coverage across all four authority levels as the deployment matures.

Data residency defines which hosting provider and geographic region stores citizen data for that country. Citizen data is always stored within the country itself where suitable infrastructure exists,

or within the EU/EEA where it does not – ensuring jurisdictional compliance independent of AGPT Ltd's UK domicile.

Localisation covers user interface language, date and number formats, currency display, and jurisdiction-specific legal text including privacy notices and terms of service adapted to local consumer protection requirements.

This configuration-driven approach enables AGPT Ltd to launch new country deployments rapidly – activating a new country requires deploying a configuration file, not writing new code – while maintaining full consistency of methodology, privacy architecture, and publication standards across all jurisdictions.

5.3.2 AGPT Ltd Responsibilities per Country

AGPT Ltd bears full and undivided responsibility for every country deployment. There are no local operators, sub-processors acting as independent controllers, or franchised entities. AGPT Ltd is the sole data controller under GDPR for all jurisdictions (§5.4.1), manages hosting infrastructure and database operations centrally (§5.3.1), builds and maintains the officials database from official registries for each country (§7.3), handles regulatory compliance with applicable national laws, and coordinates crisis response – whether legal challenge, manipulation attempt, or political pressure – from a single operational centre with country-specific adaptation as needed. This centralised model eliminates ambiguity about who is accountable when something goes wrong: the answer, in every country, is AGPT Ltd.

5.3.3 Country Adaptation and Consistency

The configuration-driven architecture described in §5.3.1 defines the boundary between what adapts and what does not. Identity verification method, official database scope, data residency, localisation, and jurisdiction-specific legal text are configurable per country. Core methodology, privacy architecture, anti-manipulation systems, publication thresholds, and ethical commitments are fixed at the platform level and cannot be overridden by any country configuration. All deployments share the same codebase – a country launch is an activation, not a fork.

5.3.4 Country Deployment Lifecycle

Each country deployment progresses through defined lifecycle stages – from waitlist through active operation – controlled centrally by AGPT Ltd. The activation criteria and progressive rollout logic are described in §7.2. If a country deployment must be suspended due to legal, regulatory, or operational reasons, or permanently decommissioned, the contingency mechanisms described in §5.7 govern the process, including GDPR-compliant data portability and erasure. Lifecycle transitions in one country do not affect the operation of any other deployment.

5.4 Data Governance and Compliance

Privacy protection is foundational to the Platform's mission. Citizens judging officials must trust that their judgments remain private, protected from governmental surveillance, political profiling, and

commercial exploitation. Without strong privacy guarantees, citizens in sensitive political contexts – polarised environments, vulnerable minorities, states with weakening democratic norms – cannot safely participate. The legal framework therefore treats privacy as mission-critical, not merely as regulatory compliance.

The privacy architecture combines legal compliance (GDPR as baseline across all jurisdictions), technical enforcement (privacy by design at the infrastructure level), and operational discipline (data minimisation, access controls, breach response). This multi-layered approach ensures that privacy protections survive even if any single layer fails.

5.4.1 GDPR as Baseline Standard

AGPT Ltd applies GDPR as the baseline data protection standard across all Platform operations, regardless of whether a given country's national framework would require it. This is a deliberate architectural decision rather than a minimum compliance exercise. GDPR provides the most comprehensive and well-understood data protection framework available, and uniform application across all 27 member states eliminates the risk of inconsistent protection levels between jurisdictions.

The Platform implements GDPR's core principles through design choices specific to its architecture. Data minimisation governs every layer: identity verification confirms that a participant is a unique adult resident of the relevant country, but the Platform does not store identity documents, precise dates of birth, or exact addresses after verification is complete. Purpose limitation is enforced architecturally – data collected for identity verification cannot be accessed by the systems that process and aggregate judgments, and judgment data cannot be linked back to verification records without a technical process that no single operator can initiate. Storage policy distinguishes between current judgment state (retained as long as the account is active, since continuous monitoring requires a live signal) and judgment history (not stored, since the Platform measures current legitimacy, not historical patterns of individual behaviour).

Accountability mechanisms are staged to the Platform's maturity. Prior to launch, AGPT Ltd completes a Data Protection Impact Assessment covering all processing activities, establishes records of processing, and designates a data protection lead. Independent audit and formal DPO appointment follow as operational scale requires, in line with the security and compliance roadmap described in Section 7.7.

5.4.2 Legal Basis for Processing

The Platform processes personal data under distinct legal bases depending on the nature of the processing activity, in accordance with GDPR Article 6(1).

Account creation, judgment submission, and access to Platform services rely on contract performance (Article 6(1)(b)). When a citizen accepts the Terms of Service, the processing necessary to operate the account, record judgments, and deliver the service follows directly from that contractual relationship. This is the Platform's primary legal basis for the majority of its processing activities. Where processing extends beyond what is strictly necessary for service delivery – such as optional communications or participation in research – explicit consent (Article 6(1)(a)) is obtained separately, and can be withdrawn without affecting the citizen's account or participation rights.

Identity verification relies on legitimate interests (Article 6(1)(f)). The Platform has a documented legitimate interest in confirming that each participant is a unique adult resident of the relevant

country, since the one-citizen-one-account guarantee is essential to the integrity of every index the Platform publishes. A balancing assessment demonstrates that this interest does not override the data subject's rights: verification data is processed by a third-party provider (§3.2), the Platform receives only a confirmation result, and no identity documents are stored after the verification process is complete. The public interest in reliable democratic accountability data serves as a supporting factor in this balancing assessment, though the Platform does not rely on public interest as an independent legal basis.

5.4.3 Cross-Border Data Flows

The Platform's data architecture is designed around country-first residency. Citizen judgment data and aggregated indices for each country are stored and processed within that country's jurisdiction, using hosting infrastructure located in the relevant member state or, where that is not available, in the nearest EU jurisdiction with equivalent data protection standards. This is an architectural constraint enforced at the tenant level, not a policy preference – the system does not permit judgment data from one country's tenant to be queried or accessed by another.

Cross-border data flows arise in limited operational scenarios. AGPT Ltd, incorporated in the United Kingdom, is the sole data controller. Routine Platform operations – judgment submission, index calculation, publication – occur entirely within each country tenant and do not require data to leave the hosting jurisdiction. However, technical support, system maintenance, and backup operations may require AGPT Ltd personnel to access country-level data from the United Kingdom. These transfers are governed by Standard Contractual Clauses as the approved transfer mechanism under GDPR Chapter V.

The United Kingdom benefits from an EU adequacy decision under Article 45 GDPR, recognising UK data protection standards as essentially equivalent to those of the EU. This provides the foundational legal basis for the controller relationship between AGPT Ltd and its EU operations. AGPT Ltd monitors the status of this adequacy decision and maintains Standard Contractual Clauses as an independent safeguard, ensuring continuity of lawful data transfers regardless of any future change in adequacy status.

5.4.4 User Rights and Data Subject Access

AGPT Ltd implements the full set of data subject rights under GDPR Articles 15–22 through Platform interface controls and documented response procedures. Citizens can access and download their current judgment state and account information, correct inaccurate account data, and export their data in a machine-readable format. These functions are available directly through the user interface without requiring a formal request, though the formal channel remains available for any right that cannot be exercised through self-service.

Account deletion requires specific attention because of its interaction with aggregated indices. When a citizen requests erasure, the Platform deletes all personal data – account information, verification status, and the link between the citizen's identity and their judgments. The judgments themselves are anonymised: the aggregate index is recalculated to reflect the removal, but no record remains that could connect the deleted account to any contribution. This is true anonymisation, not pseudonymisation – once the account is deleted, re-identification is technically impossible. The Platform does not retain judgment history (§5.4.1), so erasure affects only the current aggregate state, not historical records.

Citizens may also restrict their account temporarily, pausing participation without triggering full deletion. The right to object under Article 21 applies to processing carried out under legitimate interests – specifically, identity verification (§5.4.2). An objection to verification processing is

assessed on its merits; if upheld, it results in account suspension since the Platform cannot operate an unverified account. No automated decisions with legal or similarly significant effects are made about any individual; the Platform aggregates citizen inputs into statistical indices and does not make determinations about citizens themselves.

5.4.5 Privacy by Design and by Default

The Platform implements privacy by design and by default in accordance with GDPR Article 25 – not as a compliance layer added to a functioning system, but as a foundational architectural principle. The distinction is practical: database tables that would enable profiling, judgment history tracking, or demographic correlation of opinions do not exist in the schema. API endpoints that would return individual-level data have not been built (§5.4.10). These capabilities are not restricted by access controls or policy – they are absent from the architecture entirely. The technical details of this approach are described in Section 3; the publication thresholds that prevent statistical inference of individual opinions from aggregated data are described in §5.4.8.

Privacy by default means that a new account, with no settings adjusted, operates at the maximum privacy level the Platform offers. Judgments are never publicly visible at the individual level. No social features, follower mechanisms, or activity feeds exist that could reveal participation patterns. The only output the Platform produces from citizen activity is the aggregated index – a statistical summary in which no individual is identifiable.

5.4.6 Security Measures

AGPT Ltd implements security across four layers: infrastructure (hosting environment, network protection, backup and recovery), application (secure development practices, input validation, session management), operational (personnel controls, access governance, incident response), and third-party (vendor assessment ensuring that hosting providers, identity verification partners, and payment processors meet equivalent security standards). These layers follow a defence-in-depth model – compromise of any single layer does not expose citizen data, because each layer enforces independent controls.

The security programme is staged to the Platform's maturity, progressing from pre-launch code review and vulnerability assessment through independent penetration testing, and toward formal certification as operational scale requires. The full staging roadmap – including SOC 2 and ISO 27001 milestones – is described in §7.7. Breach response procedures comply with GDPR Article 33 notification requirements: documented detection and containment protocols, assessment of risk to data subjects, and notification to the relevant supervisory authority within 72 hours where required.

5.4.7 Data We Do Not Collect (Architectural Minimisation)

Teisond implements an aggregates-only publication model. The following categories of data are never collected, stored, linked, or made available through any Platform interface, API, or export:

Personal identifiers: names, emails, phone numbers, government IDs, and social media handles are never stored beyond the one-way cryptographic hash generated at verification. The hash confirms uniqueness; the source data is discarded.

Per-user judgment histories: no longitudinal linkages are maintained across sessions or events. The Platform maintains no 'user timelines' of opinions. A citizen's pattern of judgments across officials or over time is architecturally unrecoverable.

Reasons or explanations for judgments: the binary trust/distrust mechanism accepts no free-text justifications, tags, or categorical motives tied to individual citizens. The Platform does not ask why – only whether.

Demographic attributes and opinion correlations: age, gender, ethnicity, income, religion, education, and party affiliation are never collected. No analysis of how judgments correlate with demographic characteristics is possible, because the demographic data does not exist in the system.

Behavioural profiles: no clustering or scoring of individuals, no interest graphs, no look-alike modelling. The Platform cannot distinguish one citizen's behaviour from another's, because individual behaviour is not tracked.

Precise geolocation and device fingerprinting: no GPS-level location data, no persistent device or browser fingerprints. The Platform does not know where a citizen is when they render a judgment, nor what device they use.

Advertising and third-party trackers: no advertising pixels, retargeting beacons, or cross-site identifiers. The Platform carries no advertising and shares no data with advertising networks.

Raw personal data exports: the Platform never provides raw individual-level records to any party — including AGPT Ltd itself, researchers, media subscribers, or law enforcement – because such records do not exist in retrievable form.

These exclusions are enforced by architecture, not by policy. The data model lacks tables for individual-level opinions. Pre-aggregation gates discard any potentially identifying fields before data enters the publication pipeline. The API accepts only aggregate queries and rejects any request that could return profiling or row-level results. The user interface provides no input fields for personal details beyond the verification boundary. Necessary operational security logs – such as transient IP addresses for abuse detection – follow strict minimisation and short retention policies defined in § 5.4.6.

5.4.8 Aggregation and Publication Thresholds

To protect both individual privacy and measurement quality, the Platform publishes legitimacy indices only when they meet defined statistical thresholds. The primary safeguard is a minimum sample size – by default, no fewer than 100 independent judgments per official per publication period before any public index is displayed. This threshold ensures that published indices carry genuine statistical weight and that no small group of citizens can disproportionately determine a published number.

Beyond minimum sample size, the Platform applies k-anonymity guardrails: no metric is published if doing so risks revealing information about identifiable small groups. These guardrails are jurisdiction-configurable, allowing country deployments to set stricter thresholds where local data protection standards or demographic conditions require them – though never looser than the Platform-wide default.

Published aggregates are rounded to one decimal place to avoid false precision and reduce re-identification risk. Where an office-period combination has not yet reached the publication

threshold, the Platform displays 'Not enough judgments' rather than suppressing the entry entirely – making the absence of data visible rather than silent.

When the Platform's integrity systems detect volumetric anomalies – sudden spikes, coordinated patterns, or statistical irregularities — the affected indices remain published in their aggregated form but are flagged with a visible anomaly notice and a link to the corresponding transparency note. Suppression is avoided; disclosure is the default.

Country deployments may configure stricter local thresholds within their deployment configuration, aligned with the privacy architecture defined in § 5.4.1–5.4.6. No country deployment may set thresholds below the Platform-wide minimums.

5.4.9 Publication Policy (Public vs Subscribers vs Never)

The Platform operates three tiers of data access, distinguished not by the sensitivity of the underlying data – which is always aggregated – but by the depth of analysis and the audience it serves.

The first tier is public access, available to any visitor at no cost. This includes headline legitimacy indices for all officials who have crossed the publication threshold, confidence intervals indicating statistical reliability, trend indicators showing directional movement over time, and threshold notices where insufficient data exists. Public access also includes methodology documentation and anomaly flags – the Platform's commitment to transparency means that any detected irregularity is visible to the public, not reserved for paying subscribers.

The second tier is subscriber access, available to officials monitoring their own indices and to media, researchers, civic organisations, and consultants through paid subscriptions. Subscriber data remains strictly aggregated – no subscriber at any level receives individual-level records. What subscribers receive is analytical depth: benchmarks and comparisons across offices and institutional categories, historical time-series with configurable granularity, threshold-crossing alerts and anomaly advisories, and programmatic API access for integration with dashboards and research pipelines. The API enforces the same aggregation rules described in § 5.4.10 – it is a delivery mechanism for published aggregates, not a backdoor to underlying data.

The third tier is data that is never available – to anyone, at any price, under any circumstances. This includes individual citizen judgments, per-user judgment histories, demographic correlations of opinions, any output that could enable political profiling or re-identification of citizens, and raw datasets or row-level exports. These exclusions are not policy decisions subject to future revision – they are architectural properties of the data model. The data does not exist in a form that could be exported even if the decision were made to do so.

5.4.10 Public API and Access Controls (Aggregates-Only)

The data access architecture described in Section 4 defines three tiers of availability – public, subscriber, and restricted. The Platform's API operates strictly within this framework, serving exclusively aggregated metrics. No endpoints exist for individual-level retrieval, and no query can return data that would enable profiling or re-identification. This is not an access control decision that could be reversed – it is an architectural constraint: the endpoints that would serve individual-level data have not been built.

Query validation enforces this boundary at the infrastructure level. Requests for demographic or location breakdowns of participant opinions are rejected. Profiling queries and row-level requests

return errors regardless of the caller's authentication level. Rate limits are enforced to protect service integrity and prevent bulk extraction attempts. API keys are revocable; usage is logged at the aggregate level (never individual-level) for audit purposes.

Terms of use explicitly prohibit attempts to re-identify individuals, infer demographic correlations of opinions, or reverse-engineer individual judgments from aggregate data. Violations trigger immediate key revocation and may result in legal remedies. These terms are not aspirational – they are backed by architectural constraints that make the prohibited actions technically infeasible even if the terms were ignored.

5.4.11 Anti-Manipulation and Transparency Disclosure

The Platform's anti-manipulation framework operates at two levels: structural prevention and behavioural detection. The one-citizen-one-account guarantee, enforced through commercial identity verification (§3.2), eliminates bot-driven and duplicate-account manipulation at the architectural level. No unverified entity can participate in the judgment process. Rate-limiting mechanisms restrict how frequently a citizen can change a judgment for any given official, removing the primary vector for rapid coordinated swings. The specific parameters of these controls are not published to prevent gaming.

For coordinated campaigns by verified citizens – which are legitimate democratic activity, not fraud – the Platform employs volumetric anomaly detection, temporal clustering analysis, and campaign-burst monitoring. When suspected coordinated activity is detected, it is disclosed to users through visible labels and plain-language transparency notes rather than silently suppressed. This reflects a deliberate design choice: organised civic action, even if it produces sharp movements in an index, should not be erased without explanation. Content is removed only where it violates applicable law or Terms of Service.

The Platform applies the same visibility discipline to itself that it applies to officials. Material interventions – threshold changes, methodology updates, publication rule modifications – follow documented procedures with public changelogs. No change is made silently. Rules change only with notice, reasons, and records.

5.5 Liability and Risk Allocation

5.5.1 Defamation and Speech-Related Risks

The Platform publishes legitimacy indices that attach numerical values to named public officials. This creates inherent defamation exposure – not because the data is false, but because any publication about identifiable individuals invites legal challenge in jurisdictions where reputation is strongly protected.

Several features of the Platform's design substantially reduce this exposure. Citizen judgments are structured as binary expressions – trust or distrust – with no accompanying text, no accusations, and no factual claims about conduct. The Platform publishes only aggregated indices derived from these judgments, never individual responses. What appears on a public profile is a statistical fact: that a given percentage of participating citizens expressed trust or distrust. Reporting verifiable statistical outcomes is fundamentally different from making editorial claims about an official's character or conduct, and this distinction is recognised across both common law and civil law traditions.

The EU Digital Services Act provides the primary regulatory framework for the Platform's intermediary status. Under the DSA, platforms that host and organise user-generated content without exercising editorial control over individual contributions benefit from conditional liability protections. The Platform does not select which officials receive low or high indices, does not editorially frame the results, and does not encourage any particular judgment. It operates as infrastructure for aggregating citizen sentiment, not as a publisher advancing a position.

Defamation law varies significantly across EU member states. Common law jurisdictions (Ireland) and civil law jurisdictions (France, Germany, Poland) apply different standards regarding burden of proof, the distinction between fact and opinion, and the weight given to public interest defences. Some member states offer robust public figure doctrines; others protect official reputation more strongly. AGPT Ltd does not assume uniform legal protection across all 27 countries. National legal review is conducted for each jurisdiction prior to activation, identifying specific exposure points and adapting Terms of Service accordingly.

Despite these structural protections, residual risk is not zero. Officials may initiate strategic lawsuits against public participation (SLAPP) – actions intended not to succeed on merits but to impose financial and operational burden. The EU Anti-SLAPP Directive, currently being transposed into national law across member states, will provide procedural safeguards including early dismissal mechanisms for abusive claims. Until full transposition is complete, the Platform's mitigation rests on three pillars: dedicated media and defamation insurance coverage (§5.5.3), Terms of Service that clearly define the nature of published data as aggregated statistical outputs, and qualified legal counsel in jurisdictions identified as higher-risk during national legal review.

5.5.2 AGPT Ltd Liability Framework

As sole operator of all national Platform deployments, AGPT Ltd bears comprehensive liability for the Platform's operations across all 27 EU jurisdictions. This centralised liability model – a direct consequence of the single-entity architecture – means that citizens and subscribers in any country deal with one responsible party, governed by one set of standards, regardless of which national deployment they use.

AGPT Ltd assumes full responsibility for all operational decisions and Platform-generated content across all jurisdictions, including the accuracy of published indices within the bounds of the stated methodology. The company bears liability for all aspects of user data processing and privacy compliance under GDPR and applicable national data protection laws. Any defamation or speech-related claims arising from Platform publication – including legitimacy indices, anomaly flags, and transparency notes – are the responsibility of AGPT Ltd as publisher, not of the citizens whose aggregated judgments produced the data. Customer disputes, subscription issues, and employment claims relating to AGPT Ltd staff fall under the company's standard commercial liability.

Certain categories of harm fall outside any party's liability by their nature. Third-party attacks on Platform infrastructure – including hacking attempts and distributed denial-of-service attacks – do not create liability for either AGPT Ltd or its users, unless the damage results from AGPT Ltd's negligence in maintaining adequate security measures. Government actions compelling data disclosure or Platform shutdown in a specific jurisdiction are treated as regulatory force majeure: AGPT Ltd complies with lawful orders while exhausting available legal remedies, but does not assume liability for consequences of sovereign state action. Natural disasters, wars, and other force majeure events are governed by standard contractual provisions.

Individual users bear liability only for their own actions that violate the Platform's terms of service – such as attempting to create multiple accounts, circumventing verification controls, or misrepresenting the Platform's methodology or outputs to third parties.

AGPT Ltd maintains comprehensive liability insurance and legal reserves appropriate to the scale and jurisdictional complexity of operating civic infrastructure across multiple EU member states. Insurance coverage is reviewed annually and adjusted as the number of active country deployments and the subscriber base grow.

5.5.3 Insurance Requirements

Operating a civic accountability platform across multiple EU jurisdictions creates a specific risk profile that requires dedicated insurance coverage. The Platform publishes legitimacy indices about named public officials, creating defamation exposure. It processes citizen identity verification data, creating cyber and privacy risk. And it operates under the regulatory frameworks of multiple member states simultaneously, creating multi-jurisdictional compliance exposure.

The insurance programme will cover four domains prior to launch. Professional Liability (Errors and Omissions) addresses claims arising from incorrect index publication, methodology errors, or negligent provision of data services. Cyber Liability covers data breaches, privacy violations, cyber extortion, and business interruption resulting from attacks. Media and Defamation Liability provides specific protection for claims related to the publication of legitimacy data about identifiable officials – a risk category distinct from general professional liability. Directors and Officers coverage protects AGPT Ltd leadership from personal liability arising from organisational decisions.

Coverage levels will be established in consultation with specialist brokers experienced in technology platform and media liability, and reviewed annually as the Platform expands to additional jurisdictions. Each new jurisdiction may introduce specific liability exposures that require adjustment to the programme.

5.5.4 Tax Structure and VAT Compliance

AGPT Ltd operates as a single legal entity with a deliberately simple tax structure. The company is subject to UK Corporation Tax on its worldwide profits, filed annually with HMRC. This simplicity is an intentional choice at the current stage – a single-entity model eliminates the transfer pricing documentation, inter-company agreements, and multi-jurisdictional corporate tax filings that characterise more complex structures.

Consumer-facing VAT obligations across the 27 EU member states are managed through the Platform's Merchant of Record payment partner, which acts as the legal seller of subscription services to end users. The Merchant of Record calculates, collects, and remits VAT at the correct rate in each jurisdiction, files all required VAT returns, and issues locally compliant invoices to subscribers. AGPT Ltd receives net payouts after VAT has been handled – removing the need for direct VAT registration in individual EU member states at this stage.

Should the Platform's scale and operational requirements justify establishing a separate EU operational entity in the future – as outlined in § 7.8 – cross-border tax considerations including withholding taxes on IP licensing fees, transfer pricing compliance, and dual-jurisdiction corporate tax optimisation will be addressed at that point, with structures designed to satisfy arm's-length requirements and documented in consultation with qualified tax advisors in each relevant jurisdiction. Until then, the single-entity model keeps administrative overhead low and ensures that resources are directed toward Platform development and country activation rather than corporate structure maintenance.

5.6 Dispute Resolution

5.6.1 Governing Law and Jurisdiction

All subscriber agreements and Platform terms of service are governed by English law – a jurisdiction chosen for its well-developed body of technology and platform case law, predictable contract enforcement, and global recognition. Disputes are subject to the exclusive jurisdiction of the English courts, with one important exception: citizens and subscribers who qualify as consumers under their national law retain the right to bring claims in their local courts. Nothing in the Platform's terms limits statutory consumer rights under applicable EU or national consumer protection legislation.

For disputes involving personal data rights, users may additionally file complaints with their national data protection authority, which has independent investigative and enforcement powers under GDPR –regardless of what the Platform's terms say about jurisdiction. This dual-track protection – contractual remedies through English courts, regulatory remedies through national DPAs – ensures that no user is left without accessible recourse.

5.6.2 Escalation Process

Disputes are resolved through a staged process designed to address issues at the lowest appropriate level before formal proceedings become necessary.

The first stage is direct resolution: the user contacts AGPT Ltd's support team, which handles the matter through standard operational channels. The majority of disputes – account access issues, data display questions, subscription billing queries, and service quality concerns – are resolved at this stage.

The second stage is management review: if direct resolution fails or the dispute involves policy interpretation, account termination, or data deletion decisions, AGPT Ltd management reviews the matter and issues a formal response. Users receive written reasoning for any decision that affects their account or data.

The third stage is mediation: if management-level resolution is unsatisfactory, either party may propose mediation through a neutral third party. Mediation is non-binding but provides a structured opportunity for both sides to present their positions before an independent mediator.

The fourth stage is formal proceedings: if mediation fails, the dispute proceeds to the English courts – or, where applicable consumer protection law provides, to the user's local courts. Users may simultaneously or alternatively pursue complaints through their national data protection authority, whose powers operate independently of any contractual dispute resolution mechanism.

AGPT Ltd maintains documented complaint-handling procedures across all country deployments, ensuring consistent process, fair treatment, and timely responses regardless of which national Platform a user registered through.

5.7 Operational Resilience and Contingency

5.7.1 Contingency Architecture

The Platform's design incorporates contingency at the architectural level – not as a separate disaster recovery layer, but as a property of the choices already made. A single-entity structure with cloud-based infrastructure, a Merchant of Record payment model, and configuration-driven country deployments means that migration of any component does not require rebuilding the system – it requires changing a configuration.

Hosting infrastructure runs on a cloud provider with EU data centres. If the provider becomes unsuitable – whether due to commercial terms, regulatory pressure, or service quality – citizen data can be migrated to an alternative EU-based provider. The Platform stores no data in proprietary formats that would create lock-in; standard database exports and documented deployment procedures ensure that migration is an operational task, not an engineering project. At the current stage, migration would take days rather than hours, but the architectural prerequisites for rapid migration are in place.

Payment processing is handled by a Merchant of Record partner that manages all subscriber billing, VAT compliance, and invoicing. If the payment partner becomes unavailable, the Platform can transition to an alternative MoR provider or to direct payment processing through a conventional gateway. Subscriber relationships are maintained by the Platform – not by the payment partner – so no customer data is lost in transition.

Domain registration for teisond.com and all national subdomains is managed through standard registrars with DNS served by a global CDN provider. DNS-level failover to alternative configurations is achievable within hours.

The UK jurisdiction for AGPT Ltd was chosen for its legal stability, IP protection, and treaty network. Should conditions in the UK deteriorate materially – regulatory changes incompatible with the Platform's privacy architecture, loss of EU adequacy status, or political interference with Platform operations – AGPT Ltd has the option to re-domicile to a pre-identified EU jurisdiction. Ireland, Estonia, and the Netherlands have been assessed as viable alternatives based on their legal frameworks, technology sector infrastructure, and data protection environments. Re-domiciliation is a significant undertaking that would be executed only under serious and sustained adverse conditions, not as a routine response to regulatory inconvenience.

5.7.2 Trigger Conditions

Contingency measures are activated when conditions in any jurisdiction cross defined thresholds that threaten the Platform's core commitments – specifically, its ability to protect citizen privacy, maintain methodological independence, and publish indices without interference.

Regulatory triggers include new legislation requiring disclosure of individual citizen judgments to government authorities, prohibition of end-to-end encryption or one-way hashing for identity data, mandated content moderation incompatible with the Platform's deterministic publication methodology, or compliance burdens that make continued operation in a jurisdiction financially unviable.

Commercial triggers include termination of critical service relationships – hosting, payment processing, or identity verification – due to political pressure rather than commercial terms, or coordinated refusal of service by multiple providers in a jurisdiction.

Political triggers include direct governmental demands for preferential treatment of specific officials, use of regulatory or prosecutorial authority to pressure the Platform into altering indices or suppressing publication, or legal actions designed to harass rather than to remedy legitimate grievances.

The appropriate response to each trigger is proportionate: a regulatory change in one country may require adjusting that country's deployment configuration, while a fundamental change in UK law might necessitate corporate re-domiciliation. The contingency framework distinguishes between country-level responses – which affect a single national deployment – and entity-level responses – which affect AGPT Ltd's corporate structure. Country-level contingencies are executable by the current team; entity-level contingencies would be undertaken with professional legal and tax advisory support.

5.7.3 Preparedness and Maturity

At the current stage, contingency preparedness consists of architectural choices that preserve optionality: no vendor lock-in, no proprietary data formats, no single points of failure in the deployment pipeline, and pre-identification of alternative jurisdictions and service providers. Formal migration playbooks — documented step-by-step procedures with tested timelines for each contingency scenario – will be developed as the Platform scales and operational resources permit. This is an area where post-funding investment in operational resilience will produce concrete deliverables, as outlined in § 7.7.

5.8 Verification Independence

Civic accountability infrastructure faces a unique architectural risk: dependence on governmental systems for the identity verification that enables citizen participation. If Platform launch requires permission from the officials the Platform monitors, the resulting vulnerability is not merely operational but existential – it inverts the accountability relationship the Platform exists to create.

Teisond addresses this through verification independence as a design principle. The Platform launches with commercial identity verification providers (document check + biometric liveness) that operate across all EU member states without governmental approval or integration. These providers verify government-issued documents – passports, national ID cards – with liveness confirmation, delivering a one-person-one-account guarantee sufficient for the Platform's integrity requirements.

National eID systems under the eIDAS framework represent an upgrade path, not a precondition. Where a country's eID infrastructure is available and the responsible authority grants access, the Platform integrates it as an additional verification method – raising the assurance level from document-based to state-confirmed identity, and reducing per-verification costs. Citizens who initially verified through commercial IDV may re-verify through eID when it becomes available, with seamless account continuity.

This architecture ensures that no governmental body – whether through deliberate obstruction, bureaucratic delay, or political pressure – can prevent the Platform from operating in any EU member state. The decision to launch is the Platform's; the decision to upgrade verification is the government's.

5.9 Legal Structure Evolution

The current single-entity structure – AGPT Ltd (UK) operating all national deployments through a centralised multi-tenant architecture – is an intentional choice for the Platform's current stage. It minimises administrative overhead, eliminates inter-company complexity, and ensures that all resources are directed toward Platform development and country activation rather than corporate structure maintenance.

This structure is not permanent by design. As the Platform scales across the EU and operational complexity grows, the architecture preserves the option to establish a separate EU operational entity –likely in a jurisdiction with strong technology sector infrastructure and favourable data protection environment. Such an entity would operate the Platform's EU deployments under licence from AGPT Ltd, which would retain ownership of the methodology, software, and brand. The separation would simplify regulatory relationships within the EU, eliminate the need for an Article 27 Representative, and provide EU-domiciled institutional partners with a locally incorporated counterparty. This step is justified only when the operational and commercial benefits clearly exceed the costs of maintaining a dual-entity structure – a threshold that is assessed continuously, not scheduled to a predetermined date.

The Platform's long-term governance vision extends beyond conventional corporate structures. The founding commitment is that Teisond ultimately transitions from founder-operated infrastructure to citizen-governed civic institution – where the people who use the Platform hold meaningful governance rights over its direction. The full civic ownership architecture — including distributed infrastructure, token model, governance distribution, and the pathway from community formation to citizen ownership — is described in Section 9.

5.10 Conclusion: Legal Framework as Mission Enabler

The legal framework described in this section is designed to do one thing well: enable a UK-registered company to operate civic accountability infrastructure across 27 EU member states while protecting citizen privacy, maintaining methodological independence, and complying with the regulatory requirements of every jurisdiction in which it operates.

The framework is deliberately simple at this stage. A single legal entity serves as data controller, platform operator, and commercial counterparty across all jurisdictions. Privacy is enforced by architecture – one-way hashing, aggregates-only publication, jurisdictional data isolation – not by policy commitments that could be reversed. Liability is centralised and insured. Dispute resolution is accessible to citizens in their own jurisdictions. Contingency mechanisms preserve operational flexibility without over-engineering for scenarios that may never materialise.

As the Platform scales, the legal structure will evolve – from single-entity to dual-entity if EU operational presence justifies it, from bootstrapped contingency planning to formal resilience infrastructure, and ultimately toward governance models that give citizens a direct stake in the infrastructure they use. The framework described here is built to accommodate that evolution without requiring reconstruction – each component can be extended, replaced, or supplemented as operational reality demands.

The measure of this legal framework is not its sophistication but its adequacy: whether it enables the Platform to launch, operate, and grow while keeping every commitment made in the Neutrality

Charter and the Sovereignty and Trust Framework. The commitments are permanent. The structures that implement them will adapt.

SECTION 6: GOVERNANCE AND ETHICS

Teisond is democratic infrastructure requiring governance that ensures it serves the public interest, resists capture, and maintains integrity under pressure. Governance for a civic accountability platform carries a specific burden: the Platform applies visibility discipline to officials, and must therefore apply the same discipline to itself. This section establishes the governance framework guiding all Platform decisions – principles embedded in operating standards, enforced through architectural constraints, and subject to external accountability.

6.1 Governing Principles

6.1.1 Independence and Neutrality

The Platform maintains strict independence from governmental, partisan, and commercial interests. This independence is not aspirational – it is structurally enforced through architectural, legal, and operational commitments.

AGPT Ltd operates all national deployments as a single entity with no governmental ownership, no partisan investors, and no political affiliations. Officials being monitored cannot influence Platform operations, methodology, or data presentation – the centralised multi-tenant architecture means no local actor has access to modify code, methodology, or publication rules. The Platform monitors officials across all parties and ideologies without favour – legitimacy indices apply uniformly regardless of political category, and the Neutrality Charter, published on every national Platform, commits Teisond to this principle publicly.

While the Platform operates as a sustainable business, commercial imperatives never override the accountability mission. Subscribing officials cannot purchase preferential treatment, suppression of unfavourable indices, or influence over methodology. Calculation methods and privacy protections remain consistent regardless of external pressure, with changes occurring only through documented, transparent processes with public changelogs.

6.1.2 Transparency and Auditability

A Platform that demands transparency from public officials must practise transparency itself. Teisond implements this principle across every layer of its operations.

All calculation methods, aggregation formulas, and data processing procedures are fully documented and publicly available. The Platform's methodology is described in sufficient detail – in this White Paper and in the published technical documentation – to enable independent verification by any researcher, journalist, or interested party. As the Platform matures, additional transparency measures are planned: publication of core calculation logic as open source, third-party security audits by independent cybersecurity firms with results published in full, and annual financial

summaries showing revenue sources and major expenses. The timeline for these measures is defined in the Security and Compliance Roadmap (§ 7.7).

6.1.3 Privacy Protection as Foundational Commitment

Privacy protection is architectural, not policy-based. The system is designed so that individual judgments cannot be publicly linked to specific citizens – not through access controls that could be overridden, but through database schema that excludes the fields necessary for such linkage (see § 5.4 for comprehensive data governance framework).

The Platform never stores judgment histories, demographic correlations, or behavioural patterns. API endpoints technically cannot return individual-level data. Political profiling is architecturally impossible. This "privacy by impossibility" standard exceeds regulatory requirements and provides the strongest possible protection for citizens exercising political judgment. The question is not whether the Platform will protect privacy, but whether anyone – including the Platform's own operators – could violate it. The architectural answer is no.

6.1.4 Accessibility and Inclusion

Legitimacy monitoring infrastructure is only as strong as the breadth of participation it enables. If the Platform is accessible only to the technically literate, the urban, or the affluent, the indices it produces will reflect a skewed population – undermining both their statistical validity and their democratic legitimacy.

The Platform is designed for maximum participation: free citizen access (the judgment function is never paywalled), multilingual interfaces in every national deployment language, accessibility compliance with WCAG standards, progressive web application design eliminating app store installation barriers, and low-bandwidth optimisation for areas with limited connectivity. Accessibility is not an add-on feature – it is a design constraint applied from initial architecture, tested across devices and connection conditions, and monitored as a key performance indicator.

6.1.5 Sustained Operation and Long-Term Commitment

Legitimacy monitoring infrastructure must operate continuously and indefinitely. A Platform that collects years of citizen judgments and then shuts down does more damage than one that never existed – it betrays the trust of participants and discredits the concept.

The Platform is designed for financial sustainability from the outset – diversified revenue streams from official subscriptions and Data Access subscriptions reduce dependence on any single source. As revenue grows, the Platform will build operational reserves as a buffer against temporary disruption. Contingency planning – including hosting migration procedures, payment processing alternatives, and domain failover – is embedded in the Platform's architecture from launch (see § 5.7). AGPT Ltd's long-term commitment to sustained operation will be formalised in the Platform's Terms of Service, establishing clear obligations regarding continuity of service, advance notice of any country deployment closure, and GDPR-compliant data portability and erasure procedures for affected citizens.

6.2 Roles and Jurisdictions Passport

6.2.1 Parties and Corporate Roles

AGPT Ltd (UK) is the sole operator of the Teisond Platform across all jurisdictions. It owns the methodology, software, and brand. It operates all national deployments through a centralised multi-tenant architecture. It acts as data Controller for each national Platform under GDPR. It manages all subscription revenue, hosting infrastructure, and regulatory compliance.

National deployments are country-specific configurations of the single Teisond Platform, operated by AGPT Ltd. Each deployment stores citizen data in jurisdiction-appropriate infrastructure within the EU. National deployments are not separate legal entities – they are isolated tenants within a single codebase managed by AGPT Ltd. This structure ensures methodological consistency, privacy uniformity, and centralised incident response across all jurisdictions.

Citizens are free Platform users who verify their identity through the Platform's verification process, express trust or distrust judgments toward officials, and view published indices. Officials are monitored subjects – public figures exercising governmental authority – and optional subscribers to self-monitoring analytics services.

6.2.2 Contracting and Dispute Resolution

Subscriber Terms govern the relationship between AGPT Ltd and paying subscribers (officials, media organisations, researchers, civic organisations and consultancies). These terms are governed by English law, with local consumer protection applying where mandatory.

User Terms govern the relationship between AGPT Ltd and citizens. These terms are governed by the local law of the citizen's country of residence, with disputes subject to local courts. This ensures that citizens are never required to litigate in a foreign jurisdiction to exercise their rights.

The separation between Subscriber Terms and User Terms reflects the Platform's dual nature: a commercial product for subscribers and a civic infrastructure for citizens. The legal frameworks for each are designed to serve their respective audiences appropriately.

6.2.3 Data Protection Compliance

AGPT Ltd operates as data Controller for all national Platforms, bearing full responsibility for lawful data processing in each jurisdiction. GDPR serves as the baseline standard applied uniformly across all deployments – including any future deployments outside the EU where equivalent or higher standards are adopted. This baseline-up approach means the Platform never operates below GDPR-level protections, regardless of what a specific jurisdiction might permit.

Citizen data for each country is stored in jurisdiction-appropriate infrastructure within the EU or EEA, ensuring data residency compliance independent of AGPT Ltd's UK domicile. Cloud infrastructure providers act as Sub-processors under data processing agreements that specify technical and organisational measures, breach notification timelines, data deletion and return procedures, and audit rights. AGPT Ltd conducts due diligence on Sub-processors before engagement and maintains contractual rights to audit compliance.

As a UK-registered entity processing personal data of EU residents, AGPT Ltd has appointed an EU Representative pursuant to Article 27 GDPR. The Representative acts as the designated point of contact for data protection authorities and data subjects across all 27 member states. The Representative's identity and contact details are published in the Privacy Policy on every national Platform. This appointment supplements – but does not replace – AGPT Ltd's direct responsibility as data Controller.

6.2.4 Revenue and Taxation

AGPT Ltd collects and manages all subscription revenue centrally. The company's tax structure – UK Corporation Tax on worldwide profits, consumer-facing VAT managed through the Platform's Merchant of Record payment partner – is described in § 5.5.4. This centralised approach eliminates the inter-entity transactions, transfer pricing obligations, and multi-jurisdictional corporate filings that would arise from a more complex corporate structure.

6.2.5 Regulatory Posture

The Platform operates as an information service – it publishes aggregated statistical data about public acceptance of governmental authority. It is not a financial instrument, not a political organisation, and not a media publisher. This classification is significant because it determines which regulatory frameworks apply. The Platform does not make editorial decisions about content (indices are mathematically computed, not curated), does not offer financial products or investment advice, and does not engage in political campaigning or advocacy.

6.3 Ethical Commitments and Boundaries

6.3.1 Commitment to Democratic Values

The Platform exists to strengthen democratic accountability – not to serve as a tool for authoritarian surveillance, partisan advantage, or commercial exploitation. This commitment has operational consequences.

Country selection requires assessment of democratic conditions: the Platform does not deploy in regimes where monitoring infrastructure could be repurposed for political persecution. The Platform monitors governmental authority regardless of political orientation – it is neither a progressive nor a conservative tool, but accountability infrastructure that applies uniformly. The Neutrality Charter, published on every national Platform, formalises these commitments in a publicly binding document.

6.3.2 Commitment Against Harassment and Abuse

The Platform enables structured civic judgment, not harassment. The binary trust/distrust mechanism without free-text commentary is a deliberate design choice: it prevents the Platform from becoming a vehicle for personal attacks, defamatory statements, or coordinated abuse campaigns. There are no comment sections, no public forums, no messaging features.

Officials facing declining indices may experience discomfort – this is accountability, not abuse. The Platform clearly distinguishes between legitimate accountability pressure (public visibility of aggregated civic sentiment) and harassment (targeted personal attacks). For officials who experience threats or intimidation as a consequence of their public indices, the Platform maintains documented protocols for cooperation with law enforcement while preserving the integrity and publication of the indices themselves.

6.3.3 Commitment to Evidence-Based Operation

Methodology decisions follow evidence and statistical best practice, not political convenience or commercial pressure. Anti-manipulation algorithms are validated against empirical data before deployment. Publication thresholds are statistically grounded – minimum sample sizes, confidence intervals, and rounding rules are derived from established statistical methodology, not arbitrary choices.

Claims about Platform impact are verified through independent research rather than marketing assertions. The Platform welcomes external scrutiny of its methodology and publishes sufficient documentation for independent replication. If research reveals methodological weaknesses, the Platform commits to addressing them transparently – including publishing the finding, the response, and the timeline for remediation.

6.3.4 Commitment to User Agency and Control

Citizens control their participation at every stage. They can judge any official in the system, change their judgment as circumstances evolve, withdraw their judgment entirely, or delete their account permanently. No citizen is locked into the Platform without the ability to modify or exit their engagement – and account deletion is irreversible: all personal data is permanently removed, with only anonymised contributions to aggregate indices retained.

Officials exercise their right to respond through public statements linked to their profiles – ensuring that accountability operates reciprocally. An official facing a declining index can publish a statement visible to every citizen viewing that index, providing context, explanation, or rebuttal. The Platform facilitates this response; it does not moderate or edit it.

6.4 Stakeholder Relationships and Accountability

6.4.1 Citizens (Participating Users)

The Platform owes citizens: free access to the judgment function without paywall or conditions beyond identity verification; privacy protection through architectural impossibility of identification; transparent methodology enabling citizens to understand how their judgments become indices; responsive support for account, verification, and data access issues; and honest communication about the Platform's capabilities and limitations – including what the index does and does not measure.

Citizens owe the Platform: honest judgments reflecting genuine personal assessment (not paid-for or coerced submissions); compliance with one-person-one-account verification; and respect for the Platform's structured binary format. The Platform has no mechanism to verify sincerity – but the

one-citizen-one-account architecture ensures that even if individual judgments are strategic rather than sincere, they cannot be multiplied or amplified beyond one signal per citizen.

6.4.2 Officials (Monitored Subjects)

Officials occupy a dual position: subjects of monitoring and potential subscribers. The Platform owes officials: fair and uniform methodology applied identically regardless of party, ideology, or subscription status; privacy protection for their personal data (distinct from their public legitimacy indices); Right to Respond through published statements; access to methodology documentation sufficient to understand how indices are calculated; and equal treatment regardless of whether they subscribe.

Subscription purchases analytics access – never influence over indices, methodology, or treatment. Non-subscribing officials receive identical public indices as subscribers. This separation is absolute and architecturally enforced.

The Platform does not seek officials' consent for monitoring. Public officials exercising governmental authority are subject to civic scrutiny as an inherent aspect of their public role. This parallels media coverage, freedom of information requests, and electoral accountability – mechanisms that operate without requiring officials' agreement to be monitored.

6.4.3 Media Organisations

The Platform provides media with structured data, transparent methodology, and API access enabling integration into journalistic workflows. For media, the Platform is a permanent, verified data source that supplements (not replaces) traditional political reporting.

Media organisations maintain full editorial independence in interpreting and presenting indices. The Platform does not control media narratives, approve coverage, require favourable treatment as a condition for data access, or restrict critical reporting. If a journalist uses the data to argue that the Platform's methodology is flawed, the Platform's response is to engage with the criticism – not to restrict access.

6.4.4 Researchers and Academia

The Platform provides researchers with comprehensive data access (aggregated, anonymised datasets), transparent methodology documentation, and active cooperation in validation studies. The Platform anticipates that Public Legitimacy Analytics will become an academic subfield and actively supports its development.

Research findings – including critical ones – are welcomed as contributions to Platform improvement. Academic independence is unconditionally protected: the Platform does not review, approve, restrict, or delay research publications. If peer-reviewed research identifies methodological weaknesses, the Platform treats this as a quality signal, not a threat.

6.4.5 Civic Organisations and Civil Society

The Platform engages civil society as partners, critics, and advocates. Civic organisations can use legitimacy data for advocacy (citing indices in campaigns for governmental reform), analysis (longitudinal tracking of accountability patterns), and civic education (teaching citizens about democratic oversight mechanisms).

The Platform maintains independence from any particular civil society organisation – it does not endorse specific advocacy positions, campaigns, or political outcomes. It supports the civic ecosystem broadly by providing reliable, neutral data that any organisation can use according to its own mission.

6.4.6 Consultancies and Institutional Subscribers

The Platform provides consultancies and institutional subscribers with the same data available to media and researchers – legitimacy indices, comparative benchmarks, anomaly reports, and API access. Use cases include political advisory, stakeholder mapping, risk assessment, and strategic communications. The Platform does not endorse, validate, or take responsibility for the conclusions that consultancies draw from its data. Legitimacy indices are inputs to analysis, not conclusions – and the Platform maintains strict neutrality with respect to how subscribers apply the data in their professional practice.

6.5 Ethical Dilemmas and Resolution Frameworks

Governance principles provide compass bearings; real-world situations require judgment. This section describes four tensions the Platform will inevitably face and the frameworks for navigating them.

6.5.1 Privacy Versus Transparency

When research requests require data granularity approaching individual identification, privacy takes precedence – even at the cost of reduced analytical utility. The Platform errs systematically toward privacy protection. This is not a case-by-case judgment but a standing rule: if there is doubt about whether a data release could enable re-identification, the data is not released.

The deferred demographic extension illustrates this principle in practice. Demographic breakdowns of legitimacy data (by age, gender, region) would be analytically valuable for researchers and media. The Platform deliberately excludes this capability from the core product – it will be considered only if and when privacy safeguards are proven adequate through independent assessment. Potentially valuable research capability is deferred rather than risk citizen exposure.

6.5.2 Neutrality Versus Justice

If indices in certain regions reflect discriminatory patterns – for example, if officials from minority backgrounds consistently receive lower indices than peers with comparable conduct – the Platform faces a tension between methodological neutrality and social justice.

The Platform's resolution is clear: it measures accurately without manipulating indices to correct suspected bias. Methodology integrity is maintained unconditionally. The Platform does not adjust results to produce outcomes it considers more just – this would make the Platform an editorial actor rather than a measurement instrument, destroying the neutrality that makes the data credible.

However, neutrality does not mean silence. The Platform provides contextual information helping users interpret potential patterns. Research examining bias in legitimacy data is actively supported. But correcting injustice is society's work – through education, legislation, institutional reform – not the Platform's algorithmic adjustment.

6.5.3 Sustainability Versus Mission

When funding relationships would compromise independence, mission takes precedence regardless of financial consequences. This is not a hypothetical – it is a predictable scenario. A government may offer partnership funding contingent on methodology modifications. A commercial entity may offer sponsorship contingent on data access beyond published parameters. A political actor may offer support contingent on favourable treatment.

The Platform's resolution: revenue diversification reduces dependence on any single source, making it possible to refuse compromising offers. If diversification proves insufficient, the Platform is willing to scale back operations – reducing geographic scope, delaying expansion, or operating at reduced capacity – rather than accept funding that would compromise independence.

6.5.4 Official Responsiveness Versus Populism

Officials may abandon sound but unpopular policies to improve their indices. This is a legitimate concern – and it is inherent to democracy itself, not specific to Teisond. Elections create identical incentives: officials routinely adjust positions based on polling data, focus group research, and electoral calculus.

The Platform measures legitimacy – public acceptance of authority – not governance quality. An official who makes unpopular but wise decisions may see a declining index; an official who makes popular but harmful decisions may see a rising one. Distinguishing between responsive governance and populist pandering is society's responsibility, exercised through media scrutiny, expert analysis, institutional checks, and the voters' own judgment. The Platform provides a signal; interpreting that signal requires the full ecosystem of democratic institutions.

6.6 Governance Evolution and Future Considerations

6.6.1 Long-Term Governance Direction

As the Platform matures, governance evolves beyond the founder-led model toward structures that give citizens a direct role in stewarding the infrastructure they use. The civic ownership architecture — including distributed infrastructure, token model, governance distribution between community and professional management, and the concrete pathway from community formation to citizen ownership — is described in Section 9. The scope boundary is fixed: community governance extends to publication conventions, transparency standards, labelling formats,

governance structure changes, and strategic direction; it never extends to personal data handling, index calculation methodology, operational security, or regulatory compliance (§ 9.4).

6.6.2 Advisory Board and Stakeholder Councils

As the Platform scales beyond initial country deployments, formal advisory structures will provide external expertise and diverse perspectives without binding operational authority.

Formal advisory structures provide legitimacy-enhancing consultation and diverse perspectives without binding authority. Planned structures include an Ethics Board (political philosophy and democratic theory scholars providing guidance on ethical dilemmas and boundary cases), a Methodology Council (statisticians and political scientists advising on index calculation, threshold settings, and anti-manipulation approaches), a User Council (representative citizens providing feedback on Platform design, accessibility, and communication), and a National Platform Advisory Council (coordinating standards across country deployments).

These advisory bodies complement rather than replace operational decision-making. AGPT Ltd retains final authority on operational matters – advisory input informs decisions but does not bind them. This structure ensures that the Platform benefits from diverse expertise without creating governance paralysis or capture by advisory interests.

6.6.3 Conclusion

Governance is continuous practice, not fixed design. The Platform will face pressures to compromise independence for financial gain, manipulate methodology for political advantage, and sacrifice privacy for analytical convenience. These pressures are not hypothetical – they are inevitable consequences of operating civic accountability infrastructure in politically charged environments.

Principles provide compass bearings; navigating toward them requires judgment, institutional culture, and willingness to uphold principles against pressure. The governance framework described in this section is designed to make the right decisions easier and the wrong ones harder – through architectural constraints, public commitments, external accountability, and a culture of transparency that begins with the Platform itself.

SECTION 7: ROADMAP AND IMPLEMENTATION

This Section outlines Teisond's deployment strategy and timeline. The approach reflects the Platform's centralised, multi-tenant architecture: simultaneous presence across the EU from day one, with progressive deepening of functionality. The roadmap is designed for a digital-first platform that requires no physical infrastructure per country – enabling a pace of geographic expansion that would be impossible for a traditional service organisation.

7.1 Strategy: Simultaneous EU Presence

Teisond's launch strategy is based on a simple insight: in a digital-first platform requiring no physical infrastructure per country, there is no reason to launch sequentially. The Platform deploys national landing pages across all 27 EU member states from the first day, establishing presence, collecting early registrations, and signalling commitment to pan-European scope.

This "digital flag" approach serves three strategic purposes. First, it occupies the nascent civic trust measurement space before potential competitors – establishing Teisond as the reference platform for public legitimacy analytics across the EU. Second, it provides real-time market intelligence on where citizen demand is strongest – waitlist registration data from all 27 countries informs wave prioritisation and resource allocation. Third, it communicates that Teisond is a European infrastructure project, not a single-country experiment – a signal that matters to investors, media, officials, and institutional partners assessing the Platform's ambition and credibility.

7.2 Wave Launch: Progressive Activation

All 27 EU countries receive national landing pages from Phase 0. Full Platform activation proceeds in waves, determined by three convergence criteria: (1) identity verification infrastructure connected (commercial IDV as default, national eID where available), (2) AI-populated official database reaching sufficient coverage across all four authority levels, and (3) waitlist registrations indicating viable initial demand.

Priority markets by infrastructure readiness, institutional diversity, and early demand include Estonia (the EU's most advanced digital identity ecosystem), the Netherlands (near-universal commercial IDV coverage through iDIN and document-based providers), Poland and Spain (large-population proving grounds), and Germany (Western European institutional context). However, the actual composition of each wave is determined by which countries meet all three activation criteria first, not by a predetermined schedule.

Subsequent waves follow the same criteria-driven logic, progressively activating countries as verification access and operational readiness converge. The architecture imposes no limit on the number of countries per wave – if twelve countries reach readiness simultaneously, twelve countries launch simultaneously.

Country activation is configuration-driven: launching a new country requires deploying a country configuration file describing the national institutional structure, connecting the verification provider, and activating the AI-powered official database population. No code changes are required. The pace of expansion is constrained by operational readiness, not by engineering capacity.

7.3 Phased Functionality

Each country deployment progresses through defined phases. The timelines below are relative to each country's activation; subsequent countries follow the same sequence at accelerated pace as operational playbooks mature.

Phase 0 – Digital Flag (live from Q1 2026 for all 27 countries): National landing page with Platform description, Neutrality Charter, and waitlist registration. Each page is localised in the national language. Purpose: establish presence, measure demand, begin building the country's registration pipeline.

Phase 1 – Database Population (4–8 weeks before activation): The Platform's operational team populates the national database with Level 1 officials (members of parliament, ministers, heads of national agencies) using official government registries, parliamentary rosters, and publicly available institutional records. This ensures that citizens can begin rendering judgments on the highest-visibility officials from the first day of operation. Expansion to Levels 2–4 proceeds progressively after launch, with AI agents systematically identifying and cataloguing officials at regional, municipal, and local authority levels from public records, government rosters, and institutional directories.

Phase 2 – Citizen Registration and Judgment (activation day): identity verification activated, citizens can register and submit trust/distrust judgments. Legitimacy indices begin calculating as judgments accumulate. Indices are published once sample thresholds are met (default $n \geq 100$ per official); below-threshold officials display 'Not enough judgments.' The transition from Phase 1 to Phase 2 is the country's public launch.

Phase 3 – Full Operation (8–12 weeks after activation): Subscription services for officials and media activated. Comparative analytics, time-series data with hourly granularity, and anti-manipulation reporting fully operational. The Platform enters "quiet presence" mode: accumulating data, serving users, refining indices, and letting the numbers speak for themselves.

7.4 Election Sensitivity

Democratic legitimacy monitoring must be especially careful during election periods – the moment when public attention to political data is highest and the potential for misuse or misinterpretation is greatest.

The Platform implements freeze windows during election periods in each country. During a freeze: no new features are deployed, no methodology changes are made, no threshold adjustments are applied, and Platform communication is limited to factual operational notices. Existing indices continue to display and update normally – the freeze applies to Platform changes, not to citizen judgments. Citizens continue judging officials throughout the election period; the Platform simply refrains from any action that could be perceived as influencing the electoral process.

Specific freeze periods are configured per country in accordance with local electoral legislation. In countries with staggered regional elections, freezes apply to the affected region's officials while other regions continue normal operation. The freeze mechanism is documented in the Platform's public methodology and activated automatically based on the country configuration file.

7.5 Long-Term Vision

By end of 2027, Teisond aims to operate as standard civic infrastructure across the European Union – a permanent, neutral, privacy-preserving channel through which citizens in any EU member state can express and track their acceptance of officials at every level of government.

The Platform's success will be measured not by user engagement metrics but by whether legitimacy indices become a routine reference point in democratic discourse – cited by media alongside election results and economic indicators, consulted by officials as career-relevant feedback, used by researchers as a standard dataset for accountability studies, and trusted by citizens as a reliable reflection of collective judgment.

Year 2: all activated Platforms operational and approaching financial self-sustainability. Year 5: 27 EU member states fully covered; legitimacy monitoring recognised as a legitimate accountability mechanism alongside elections, courts, media, and civil society. Year 10: 40–60 democracies covering 1–2 billion population; the Platform becomes standard democratic infrastructure. Beyond Year 10: officials who grew up under continuous legitimacy monitoring govern differently from the generation that knew only episodic accountability.

The timeline is ambitious but grounded in the Platform's architecture: config-driven deployment, automation-first operations, and sublinear cost scaling mean that geographic expansion is constrained by market readiness and verification infrastructure readiness, not by engineering or operational capacity.

7.6 Infrastructure and Partner Ecosystem

The Platform's operational infrastructure relies on a curated ecosystem of external service providers, each selected for startup-stage compatibility, EU regulatory alignment, and the ability to scale alongside the Platform without requiring provider migration. The guiding principle is deliberate simplicity: each function is served by one primary provider, with architectural independence preserved to enable switching if commercial terms, service quality, or regulatory conditions change.

Identity verification is provided by a commercial IDV provider offering document scanning with biometric liveness checks across all 27 EU member states. The provider is selected for EU-native

infrastructure, broad document coverage, GDPR-compliant architecture, and white-label capability allowing the verification experience to feel native to each national Platform. Where national eID systems become available and connected, they supplement – but do not replace – the commercial verification layer, ensuring the Platform never depends on government-controlled identity infrastructure for its operation.

Payment processing and tax compliance are managed through a Merchant of Record partner. The Merchant of Record acts as the legal seller of all subscription services, handling payment collection, VAT calculation and filing across all 27 EU jurisdictions, invoice generation, chargeback management, and currency conversion. AGPT Ltd receives net payouts after all tax obligations have been settled – eliminating the need for direct VAT registration in individual EU member states and removing one of the most significant operational burdens facing any company selling digital services across the EU.

Cloud infrastructure is hosted by an EU-based provider with data centres in the European Union, ensuring jurisdictional data residency compliance. The Platform uses standard database formats and documented deployment procedures – no proprietary data formats, no vendor lock-in – preserving the ability to migrate to an alternative provider if necessary.

EU GDPR representation is provided by a specialist compliance firm appointed pursuant to Article 27 GDPR, serving as the designated point of contact for data protection authorities and data subjects across all 27 member states. The representative also provides privacy operations tooling – including a system for managing data subject requests and an incident management system for data breach notification – giving the Platform baseline compliance infrastructure from day one.

Legal services are provided through a credit-based legal platform specialising in cross-border startup operations. This model provides access to GDPR specialists, contract lawyers, and corporate counsel across multiple jurisdictions without retainer commitments – legal capacity scales with operational needs rather than imposing fixed costs during pre-revenue periods. For matters requiring specialist English law or complex regulatory expertise, a boutique digital law firm with dual qualification in English and EU law serves as secondary counsel.

Accounting and financial compliance are managed by a UK-based firm specialising exclusively in startups and scale-ups, providing bookkeeping, tax filing, payroll, and investor-ready financial reporting. The choice of a startup-specialist accountant – rather than a general-purpose provider – reflects the expectation that the company's financial reporting needs will evolve rapidly as it transitions from pre-revenue to active operations to fundraising.

Automated communications – transactional notifications, subscription confirmations, alert emails, and data subject request acknowledgements – are delivered through a dedicated email infrastructure provider, separate from the company's person-to-person email. This separation ensures that automated Platform communications maintain high deliverability without interference from general business correspondence.

All external service providers processing personal data on behalf of AGPT Ltd operate under Data Processing Agreements specifying technical and organisational measures, breach notification timelines, data deletion procedures, and audit rights. DPA execution follows a staged schedule aligned with the Platform's deployment phases: providers handling citizen data (cloud infrastructure, identity verification) require executed DPAs before the first country activation; providers handling subscriber data (payment processing) require executed DPAs before subscription services launch; providers handling only operational data (email infrastructure, accounting) require executed DPAs before the relevant service begins processing personal data. Legal counsel reviews all DPAs for adequacy before execution.

7.7 Security and Compliance Roadmap

The Platform's security posture matures through defined stages, each corresponding to a level of operational scale and institutional exposure. The approach is pragmatic: security investment is directed where it produces the greatest risk reduction at each stage, rather than pursuing comprehensive certification before the Platform has users.

Stage 1 covers the pre-launch and initial activation period. Security at this stage is embedded in the development process itself. The Platform is built using AI-assisted development tools with a dual-agent methodology: one agent writes code, a second agent performs security-focused review – examining authentication flows, data encryption implementation, input validation, injection protection, and cross-site scripting defences. Before the first country activation, an independent security engineer conducts a focused code review of the Platform's critical paths: identity verification integration, cryptographic hashing pipeline, data isolation between country tenants, and publication controls. This review produces a documented assessment that serves as the Platform's initial security baseline.

Stage 2 begins once the first countries are active and citizens are using the Platform. At this point, the Platform commissions a formal penetration test from an independent cybersecurity firm. The penetration test simulates real-world attacks against the Platform's production infrastructure – probing for vulnerabilities that code review alone cannot detect, including configuration errors, network-level exposures, and authentication bypass scenarios. The resulting report – including all findings, severity ratings, and remediation steps – becomes part of the Platform's security documentation. This document can be shared with data protection authorities, institutional partners, and investors as evidence of the Platform's security posture.

Stage 3 corresponds to the Platform's growth phase, typically following the first funding round. As institutional subscribers – media organisations, research institutions, civic organisations – begin relying on the Platform's data, they will require evidence of systematic security management. SOC 2 Type II certification provides this evidence: an independent auditor examines the Platform's security controls, availability, processing integrity, confidentiality, and privacy over a sustained observation period (typically 6–12 months). SOC 2 certification is increasingly a prerequisite for enterprise contracts and institutional partnerships. The certification process also imposes internal discipline – formalising incident response procedures, access control policies, change management processes, and vendor management standards.

Stage 4 addresses long-term institutional credibility. ISO 27001 certification – the international standard for information security management systems – provides the most widely recognised evidence that the Platform manages information security through a comprehensive, auditable framework. This level of certification is relevant when the Platform operates at scale across multiple jurisdictions, processes significant volumes of citizen data, and engages with governmental or quasi-governmental institutional partners who require ISO 27001 as a baseline condition. The certification is also a signal to data protection authorities that the Platform takes its security obligations seriously beyond minimum regulatory compliance.

Each stage builds on the previous one: code review findings inform penetration test scope, penetration test findings inform SOC 2 control design, and SOC 2 controls form the foundation of the ISO 27001 management system. This progressive approach ensures that security investment produces cumulative value rather than redundant assessments.

7.8 Funding Strategy and Capital Allocation

Teisond's development through the pre-launch and initial activation phases is entirely self-funded by the founder. The Platform, the White Paper, the governance framework, the full site across 27 countries, and all operational infrastructure have been built without external capital – by a solo founder using AI-assisted development tools. This bootstrapped phase demonstrates execution capability and product viability before any external investment is sought.

The bootstrapped approach is not a constraint imposed by necessity – it is a deliberate strategic choice. Building the Platform to functional completion before raising capital means that investor conversations begin from a position of demonstrated execution, not speculative promises. The product exists. The documentation is published. The infrastructure is operational. The question for an investor is not whether the founder can build what is described, but whether the resources exist to scale what has already been built.

External funding – at the angel investment stage – is required for a specific, bounded set of scaling needs that cannot be addressed through operational revenue alone, because they must be in place before or concurrent with the revenue growth they enable:

Security certification and independent auditing represent the largest single investment category. A formal penetration test, SOC 2 Type II certification, and the preparatory work for ISO 27001 require professional services that cannot be bootstrapped. These certifications are prerequisites for institutional credibility – media organisations, research institutions, and enterprise partners require evidence of systematic security management before committing to data subscriptions or integration partnerships. Estimated cost: €20,000–45,000.

Establishment of an EU operational entity – in a jurisdiction such as Estonia, Ireland, or the Netherlands – eliminates the need for an Article 27 Representative, simplifies VAT and regulatory relationships within the EU, provides institutional partners with a locally incorporated counterparty, and strengthens the Platform's positioning as European civic infrastructure operated from within the EU. Estimated cost: €5,000–10,000 for establishment, approximately €3,000 per year for maintenance.

Intellectual property protection through trademark registration in the EU (EUIPO) and UK (IPO) across the relevant service classes, ensuring that the Teisond brand is protected in every jurisdiction where the Platform operates. Estimated cost: €2,000–3,000.

Legal and compliance capacity expansion – increasing the credit allocation with the Platform's legal services provider and engaging specialist counsel for investor documentation (SEIS/EIS advance assurance, subscription agreements, shareholder protections). Estimated cost: €10,000–20,000.

Operational runway – covering fixed operational costs (EU Representative, accounting, hosting, email infrastructure) for 12 months beyond the point of first revenue, providing the financial stability necessary to focus on user acquisition and country activation without premature pressure to optimise for short-term revenue. Estimated cost: €15,000–25,000.

The total funding requirement is estimated at €50,000–100,000 – a range consistent with typical angel investment rounds for early-stage European startups. Every line item addresses a specific scaling barrier between a functioning product and an institutionally credible, commercially operational platform. There are no speculative R&D costs, no team expansion beyond immediate operational needs, and no capital allocated to features that do not yet have validated demand.

This capital allocation philosophy reflects the same principle that guided the bootstrapped phase: spend on what produces measurable results, defer what can be deferred without compromising the

mission, and maintain the financial discipline that makes the Platform's long-term sustainability credible.

SECTION 8: TEAM AND ORGANISATION

8.1 Organisational Philosophy

Teisond is designed as an automation-first organisation. The Platform's multi-tenant architecture, config-driven country deployment, and algorithmic index calculation mean that routine operations – publishing indices, enforcing thresholds, detecting anomalies, managing country configurations – run without human intervention. People handle exceptions, strategy, stakeholder relationships, and quality oversight. The system does the work; the team directs and governs it.

This is not a temporary staffing compromise – it is a permanent architectural choice. Adding a new country deployment does not require hiring a country team; it requires deploying a configuration file and connecting a verification provider. Scaling from 5 to 27 countries does not proportionally increase headcount. The marginal cost of geographic expansion is infrastructure, not personnel. This means Teisond can achieve pan-European presence with a team an order of magnitude smaller than a traditional multi-country service organisation would require.

The organisational model follows from the Platform's mission. Civic accountability infrastructure must be financially sustainable without grant dependence or extractive monetisation. Automation-first operations make this possible: low operational costs relative to deployment scale enable profitability at modest adoption rates, reducing the pressure to compromise mission for revenue.

8.2 Founder and Leadership

Oleksiy Loboyko – Founder, CEO. Nearly five years developing the Teisond concept across its theoretical, methodological, legal, and architectural dimensions. Background in strategic communications, political analysis, and civic technology. Based in Ukraine; operational base transitioning to the United Kingdom upon AGPT Ltd activation.

The founder's Ukrainian background is relevant context, not a liability. Teisond was conceived by someone who has experienced firsthand what happens when accountability infrastructure is absent – where the consequences of unchecked authority are not academic abstractions but lived realities. This experience informs the Platform's design: the insistence on privacy by construction (because citizens in vulnerable environments cannot afford exposure), the architectural impossibility of political profiling (because such tools have been abused elsewhere), and the commitment to neutrality (because a tool that serves one political faction is a weapon, not infrastructure).

The founding stage is intentionally solo. The concept, methodology, legal architecture, and technical design have been developed to a level of completeness that allows recruitment to target execution rather than ideation. The leadership team expands as the Platform transitions from documentation to implementation – with recruitment prioritising mission alignment alongside technical capability.

A note on long-term intent: the founder's vision for Teisond is not a company to be permanently owned but infrastructure to be progressively shared. The legal architecture deliberately preserves pathways toward citizen co-governance (see § 6.6 and § 9.5). The founder's role is to launch and prove the model – not to retain permanent control over civic infrastructure that, by its nature, should belong to the public.

8.3 Team Development

Teisond's team grows in response to operational milestones, not predetermined timelines. The automation-first architecture means the Platform can launch and operate initial country deployments with a small core team; subsequent hiring is triggered by operational needs as they emerge, not by org-chart projections written years in advance.

The first hires address the functions that cannot be automated: technical leadership (platform architecture, security, privacy engineering), country expansion management (verification provider integration, official database sourcing, regulatory navigation), and legal and compliance oversight (multi-jurisdictional GDPR compliance, operating agreements, dispute resolution). These are the roles where human judgment, cross-jurisdictional expertise, and stakeholder relationships are irreplaceable.

As the Platform scales beyond initial waves, the team expands into data science (methodology refinement, anomaly detection improvement), product management (user experience, subscriber value), and regional coordination. The principle remains constant: hire when a function demonstrably requires a dedicated person, not when an org chart says it should exist. Every role must justify itself against the alternative of automation or outsourcing.

Remote-first operations enable global talent access while reducing overhead. Compensation balances competitiveness with social enterprise positioning – competitive base salaries supplemented by equity participation rather than above-market cash. The goal is to attract people who are drawn to the mission and capable of the work, not people optimising for compensation alone.

8.4 Advisory Board

External advisors provide expertise the core team does not possess and institutional credibility that a startup cannot generate on its own. The advisory board is structured around the domains where independent perspective matters most: political science and democratic theory (ensuring the Platform's conceptual framework withstands academic scrutiny), data protection law (navigating GDPR compliance across 27 jurisdictions), platform business models (validating commercial assumptions and pricing strategy), and civic technology (learning from predecessors' successes and failures).

Advisory board members provide guidance without operational authority – they inform decisions but do not make them. This structure ensures diverse perspectives reach the leadership team without creating governance paralysis or diffusing accountability. Advisors are selected for substantive expertise and independence; the Platform does not appoint advisors for their names or connections.

8.5 Governance and Succession

AGPT Ltd governance follows a natural trajectory: founder-led during the startup phase, transitioning to board governance as the Platform matures and the team grows. The timing of this transition is driven by operational scale and institutional readiness, not by a predetermined date.

Succession planning begins early – not because founder departure is anticipated, but because infrastructure that serves millions of citizens across 27 countries must not depend on any single individual. Documentation of methodology, strategic rationale, stakeholder relationships, and operational procedures ensures continuity. Key-person dependencies are systematically identified and mitigated through knowledge distribution, cross-training, and documented decision frameworks.

The long-term governance vision – described in § 6.6 and § 10.5 – includes potential citizen participation in Platform governance. Whether governance evolves toward decentralised participation or remains within a traditional board structure, the principle is the same: the Platform's governance must be resilient enough to survive the departure of any individual, including its founder.

8.6 Community as Institutional Foundation

Civic accountability infrastructure that monitors officials exercising governmental authority operates in an inherently adversarial environment. Legal protections (Section 5) and architectural resilience (Section 3) address institutional and technical dimensions of this exposure. Neither is sufficient alone. A Platform that relies exclusively on legal rights and server architecture remains vulnerable to a class of threats that neither lawyers nor engineers can neutralise: sustained political pressure, coordinated reputational campaigns, or strategic withdrawal of commercial partnerships under governmental influence. The missing layer is social: a publicly visible community whose documented commitment to the Platform's mission makes interference politically costly.

This community is not expected to emerge organically. It is built through three deliberate mechanisms, deployed in sequence as the Platform matures. The first mechanism is the Platform itself. From launch, every verified citizen who renders judgment becomes a documented stakeholder – a person who has invested identity verification, time, and a civic act into infrastructure that did not exist before. The accumulated base of verified participants is the Platform's most fundamental source of institutional weight: not users of a service, but citizens exercising a right they did not previously have. The second mechanism is professional stakeholder formation through the Data Access tier. Journalists who build editorial workflows around legitimacy indices, researchers who publish studies using Platform data, civic organisations that integrate benchmarks into advocacy – each develops a professional dependency on the Platform's continuity. This layer is smaller in number but disproportionate in influence: a newspaper that regularly cites NOLI scores has both the incentive and the public reach to challenge any attempt to suppress the data source. The third mechanism is targeted public mobilisation: civic crowdfunding campaigns, a public supporter registry, and academic validation partnerships with research institutions whose independent assessment of the Platform's methodology provides external credibility. These instruments convert diffuse public sympathy into a documented, countable, and citable base of support.

The strategic logic is straightforward. Pressure against a solo founder operating a UK limited company is low-cost and low-risk for the aggressor. Pressure against a Platform with thousands of

verified citizens, dozens of professional subscribers whose published work depends on its data, a public register of named supporters, and independent academic endorsement of its methodology is a qualitatively different proposition – one that generates media coverage, parliamentary questions, and reputational damage for the entity applying pressure.

Community formation does not replace legal and architectural protections – it complements them. All three layers operate simultaneously: legal structures define rights and remedies, architecture prevents single points of failure, and community raises the political cost of interference. The Platform's institutional resilience is the product of all three, not any one in isolation.

The community formation strategy described above has a direct relationship to the Platform's civic ownership trajectory (Section 9). Founding supporters who contribute to the Platform's early community – whether through civic participation, professional engagement, or financial backing via civic crowdfunding – represent the natural constituency for the governance mechanisms described in §§ 9.3–9.4. When the Platform issues utility tokens, early documented supporters constitute the first cohort with a demonstrated stake in the infrastructure they helped establish. Community formation is therefore not merely a protective strategy – it is the first practical step toward citizen ownership.

SECTION 9: CIVIC OWNERSHIP ARCHITECTURE

9.1 The Principle: Value Belongs to Those Who Create It

The Platform generates an informational product – continuous, structured, verified legitimacy data covering officials across 27 EU member states. This product has measurable economic value: media organisations, researchers, consultancies, and civic organisations pay for access to it (Section 4). Yet the raw material from which this product is constructed – individual civic judgments rendered by verified citizens – is contributed by users who, under conventional platform economics, would receive nothing in return. The value is extracted from participation; the profit accrues to the entity.

This model is familiar. Social media platforms monetise user-generated content, attention, and behavioural data without sharing ownership or revenue. Survey companies maintain respondent panels whose aggregate output drives multi-billion-dollar industries while individual panellists receive token compensation or none at all. In each case, the underlying economic logic is the same: the platform captures value created by its users and distributes it to shareholders who may never have used the product.

Teisond rejects this logic as both ethically inconsistent and architecturally unsound. A Platform that asks citizens to render civic judgment on officials exercising authority over their lives – and then treats those judgments as raw material for a privately owned business – reproduces the very asymmetry it exists to correct. Citizens would exchange one form of voicelessness for another: instead of having no mechanism to express judgment, they would have a mechanism whose output they neither own nor control.

The architectural commitment is therefore explicit: the Platform is designed from its first line of code for eventual citizen ownership. Every mechanism described in this White Paper – from identity architecture to community formation (§ 8.6) to the governance trajectory described later in this Section – is a step on this path. The destination is fixed; the specific legal and technical forms through which ownership is expressed will be determined by regulatory conditions and technological maturity at the time of implementation.

9.2 Distributed Infrastructure

Citizen ownership of civic infrastructure is meaningless if the infrastructure itself resides on servers controlled by a single entity. A company that holds all data can be compelled to delete it, alter it, or surrender it – regardless of what governance documents promise. Genuine ownership requires that the infrastructure physically cannot be captured, because it does not exist in any single location.

Teisond's long-term architectural trajectory is therefore distributed: national databases replicated across the devices of participating users, with node operators maintaining local copies of the cryptographically verified dataset for their jurisdiction. Blockchain technology serves here not as a fashionable label but as the only proven mechanism for maintaining a tamper-evident, collectively held ledger where no single party – including the founding entity – can unilaterally modify or delete records. Each judgment, once cryptographically committed, becomes part of a permanent, auditable, collectively maintained dataset.

Node operators – citizens who dedicate device storage and processing capacity to maintaining the network – occupy a structurally privileged position. They bear the cost of infrastructure maintenance that would otherwise fall on a centralised provider, and their contribution is both measurable and essential. This privileged position is reflected in governance rights: node operators receive enhanced participation in Platform governance decisions proportional to their infrastructural contribution, without violating the one-citizen-one-account principle that governs judgment itself.

The transition from centralised to distributed infrastructure is progressive, not instantaneous. At launch and through the early operational phase, the Platform operates on conventional cloud infrastructure (Section 3) – because distributed systems require a critical mass of participants to function reliably, and because the Platform must demonstrate methodological integrity before asking citizens to host its data. The migration path is designed into the architecture from the outset: data formats, cryptographic commitments, and API structures are built to support eventual distribution without requiring reconstruction.

9.3 Token Model and Ownership Pathway

Ownership in distributed civic infrastructure requires a mechanism that is verifiable, transferable, and resistant to centralised revocation. Tokenisation – the representation of rights and stakes as cryptographic tokens on a distributed ledger – provides this mechanism.

The Platform's token model proceeds in two stages, each triggered by community readiness rather than a predetermined schedule. The first stage is utility token issuance. Utility tokens grant holders specific functional rights within the Platform ecosystem: participation in governance decisions (§ 9.4), priority access to Data Access tiers, eligibility for node operator status, and recognition as documented stakeholders in the Platform's institutional base. Utility tokens do not represent equity, do not pay dividends, and do not constitute financial instruments. They are functional keys to a civic infrastructure, issued to citizens who contribute to its existence – whether through civic participation, financial support via crowdfunding campaigns (§ 8.6), or infrastructural contribution as node operators.

The second stage is conversion from utility to security tokens – a transformation that formalises citizen ownership in legal and economic terms. This conversion is not a management decision. It is triggered exclusively by a supermajority decision of token holders themselves – requiring approval of no less than seventy-five per cent of participating holders. The threshold is deliberately high: the conversion changes the fundamental legal character of the tokens and introduces regulatory obligations under the EU's Markets in Crypto-Assets (MiCA) framework, national securities regulations, and potentially cross-border compliance requirements. A decision of this magnitude must reflect overwhelming community consensus, not a narrow majority.

The regulatory landscape for both utility and security tokens in the EU is complex but navigable. MiCA, fully applicable since December 2024, establishes a comprehensive framework for crypto-asset issuance, service provision, and market conduct. The Platform's token design is informed by this framework from the outset – not as an afterthought but as a design constraint. Compliance with

MiCA's requirements for white papers (in the regulatory sense), reserve obligations, and consumer protection is incorporated into the token architecture before issuance, not retrofitted after. This is an area where professional legal and regulatory advisory is essential, and the costs of obtaining it are among the scaling barriers identified in § 7.8.

9.4 Governance Distribution

The question of what citizens govern – and what remains under professional management – is not abstract. It determines whether citizen ownership is substantive or decorative. Teisond's answer is guided by a simple test: if a decision requires specialised expertise and carries legal liability, it belongs to professional management; if it shapes the Platform's relationship with its users and the public meaning of its output, it belongs to the community.

Community governance scope includes: publication conventions and labelling standards – how indices are presented, what contextual information accompanies them, and what terminology is used in public communications. Community guidelines – the norms governing citizen conduct on the Platform. Governance structure changes – including the threshold and process for utility-to-security token conversion. Strategic direction – whether to expand to new jurisdictions, whether to develop new data products, and how to allocate community-controlled resources. These are decisions that affect the Platform's civic character and public legitimacy – precisely the domain where citizen judgment is not merely appropriate but essential.

Professional maintenance scope includes: legal compliance across all operating jurisdictions, security operations and incident response, identity verification provider relationships, index calculation methodology, personal data handling, and regulatory reporting. These functions require specialised expertise, carry personal and institutional liability, and must be executed with speed and consistency that deliberative governance cannot provide. They are analogous to the role of a management company in a condominium: residents decide whether to renovate the lobby; professionals maintain the electrical system. Residents can replace the management company; they do not rewire the building themselves.

This distribution is not static. As the community matures and its governance capacity grows, the boundary between community and professional scope may shift. The direction of this shift is always toward greater community authority – never the reverse. Functions that move into community governance do not return to professional management. The principle is ratchet, not pendulum.

9.5 From Community Formation to Citizen Ownership

The pathway from a founder-led startup to citizen-owned civic infrastructure is not a single transition but a sequence of concrete steps, each building on the institutional foundation laid by the previous one.

The first step is community formation, described in § 8.6. Verified citizens who render judgment, Data Access subscribers whose professional work depends on Platform data, and civic crowdfunding backers who financially support the Platform's expansion collectively constitute the documented stakeholder base from which all subsequent governance structures emerge. This step requires no token issuance, no blockchain infrastructure, and no regulatory approval – only a

functioning Platform and a deliberate strategy for converting users and supporters into a visible, countable community.

The second step is utility token issuance. Once the community has reached sufficient scale and the Platform has demonstrated operational stability, governance tokens are issued to existing stakeholders – verified citizens, Data Access subscribers, crowdfunding backers, and node operators. The issuance criteria and allocation methodology are published in advance and subject to community consultation before implementation. Early supporters – particularly those who backed the Platform through civic crowdfunding before tokens existed – receive allocation that reflects their contribution to the Platform's institutional foundation at a stage when its survival was not guaranteed.

The third step is infrastructure distribution. The Platform's data layer migrates progressively from centralised cloud hosting to a distributed network maintained by node operators. This migration is technically complex and operationally sensitive – it proceeds jurisdiction by jurisdiction, with extensive testing and fallback mechanisms, and only when the node operator base in a given country is sufficient to guarantee data availability and integrity.

The fourth step is the community's own decision regarding security token conversion – the moment when citizen ownership is formalised in legal and economic terms. This step is not scheduled, not assumed, and not within the founder's authority to initiate. It belongs to the community, requires supermajority approval, and may or may not occur depending on the community's assessment of regulatory conditions, institutional readiness, and collective preference. The founder's role is to ensure that the option exists – not to exercise it.

The measure of success for this trajectory is specific: a Platform that continues to operate with full methodological integrity, privacy protection, and public trust – without any single individual, company, or jurisdiction being necessary for its survival. Not a platform that belongs to its founder and serves its users, but a protocol that belongs to those who use it and is maintained by those it serves.

SECTION 10: CONCLUSION

10.1 The Case for Continuous Legitimacy Monitoring

This White Paper has established that modern representative democracy suffers from two interconnected accountability deficits: the Extra-Electoral Voter Influence Deficit affecting elected officials between elections, and the Legitimate Public Influence Loop Deficit leaving the vastly larger apparatus of appointed administrators outside any structured citizen oversight (§ 1.1). Both deficits share a single root cause: the structural absence of mutual accountability infrastructure between citizens and officials (§ 2.9).

Teisond addresses both through a single innovation: permanent public infrastructure enabling any verified citizen to express trust or distrust toward any official exercising governmental authority – continuously, privately, and with results aggregated into public indices that make collective civic judgment visible.

Legitimacy is not a status conferred once and presumed thereafter. It is a continuous relationship between those who exercise authority and those over whom it is exercised – a relationship that strengthens or erodes with every interaction, every decision, every encounter. No mechanism existed to make this relationship visible. The Platform provides one.

The Platform does not impose accountability from above. It opens an information channel between governing and governed – groups whose relationship has been structurally unidirectional. Authority flows downward continuously; the citizen's capacity to signal acceptance or rejection of that authority now flows upward through a permanent, verified, privacy-preserving mechanism. What emerges from this bilateral exchange – accountability, mutual recognition, reduced tension – is not the Platform's direct function but the social capital generated when a blocked channel is finally opened.

10.2 Why the Teisond Solution is Viable Now

Several convergent developments – technological, social, and commercial – make this viable now in a way that would have been impossible a decade ago. Commercial identity verification technology – document scanning with biometric liveness checks – now enables reliable one-citizen-one-account verification across all 27 EU member states, preventing manipulation while preserving privacy; where national eID systems are available, they provide an additional assurance layer without the Platform depending on any government-controlled infrastructure for its operation. Cloud infrastructure has matured to the point where platforms serving millions of users across dozens of jurisdictions can scale without prohibitive capital requirements. Declining institutional trust across European democracies has created genuine citizen demand for transparency tools that restore agency rather than merely document dissatisfaction. The civic technology ecosystem has produced enough successes – and enough instructive failures – to confirm that social innovations

work when their design respects both human psychology and institutional reality. And proven market demand for political analytics demonstrates that subscription-based sustainability is achievable without grant dependence or philanthropic patronage.

Most importantly: Teisond addresses a permanent condition – citizens living under governmental authority with no routine mechanism to signal whether they accept it as legitimate. This is not a gap between elections; it is a structural absence that elections themselves do not remedy. The Platform closes it.

The Platform's viability derives not from idealistic assumptions about civic virtue but from addressing specific unmet psychological needs: restoration of dignity and agency in relations with governmental authority (§ 2.5). Officials subscribe from career self-interest; media integrate indices from editorial pragmatism; researchers access data from academic motives. Each stakeholder category interacts with the Platform from their own interests, yet their collective participation generates shared democratic infrastructure. Teisond works with human nature as it exists, transforming existing motivations into an applied information product of universal application.

10.3 Stakeholder Value Proposition

Citizens gain what no existing mechanism provides: a permanent, private, low-effort channel to render civic judgment on any official who exercises authority over their lives – from a cabinet minister to a local school principal.The psychological value is immediate: restoration of agency and dignity in encounters with governmental authority where helplessness previously prevailed.

Officials gain continuous, verified feedback that no other source provides – not episodic poll snapshots commissioned by media, but a permanent signal reflecting how citizens judge their exercise of authority. Whether an official subscribes from career self-interest, competitive awareness, or genuine responsiveness, the analytical layer provides insight unavailable through any other channel. The Right to Respond mechanism ensures accountability operates reciprocally.

Media gain a permanent, structured data source that replaces ad-hoc polling commissions. Legitimacy indices provide continuous news hooks beyond episodic scandals, enable data journalism formats previously impossible at local and regional level, and offer citation-ready metrics with transparent methodology – a structural advantage in an industry under constant cost pressure.

Researchers gain a new empirical data layer for governance studies: office-period as analytical unit, standardised metrics with transparent publication rules, and cross-country comparative possibilities across 27 EU member states. Public Legitimacy Analytics as an academic subfield begins with the data the Platform provides.

Investors and strategic partners see a Blue Ocean market positioning – Public Legitimacy Analytics as a new segment within the Public Opinion Research market – with a mission-aligned revenue model, scalable multi-tenant architecture with sublinear cost growth, and first-mover advantage in a nascent civic infrastructure category.

For democracy broadly, the Platform normalises continuous civic oversight as a complement to elections, courts, media, and civil society. It provides early detection of legitimacy crises, reduces accumulated frustration that otherwise finds expression in protest or withdrawal, and supports the long-term resilience of democratic institutions by making the consent of the governed as continuously visible as the authority of those who govern.

10.4 What the Platform Does Not Provide

Honesty about limitations earns more trust than omissions. Teisond is a specific instrument with clear boundaries.

The Platform does not guarantee governmental responsiveness. Officials can ignore low indices, accepting reputational losses. The Platform creates pressure through visibility, not legal obligation. Whether officials respond to that pressure is a political and cultural question, not a technical one.

The Platform does not prevent bad governance or corruption. It measures public acceptance of authority, not governance quality – competent officials making unpopular decisions may have low indices; charismatic but incompetent officials may maintain high ones. Distinguishing between responsive governance and populism is society's responsibility, exercised through the full ecosystem of democratic institutions.

The Platform does not solve representativeness challenges. Published indices reflect participating citizens, not the entire population. Transparency about sample sizes, confidence intervals, and threshold notices helps interpret but does not eliminate participation bias. The Platform discloses these limitations alongside every published index.

The Platform does not prevent manipulation. Coordinated campaigns by real citizens can move indices through legitimate mobilisation. The Platform makes such activity visible through anomaly flags and transparency notes rather than censoring it – respecting the distinction between organised civic action and illegitimate manipulation.

The Platform does not replace other democratic mechanisms. Elections, courts, media, and civil society retain their full functions. The Platform supplements them with continuous civic judgment data – filling a specific structural gap, not claiming to be a comprehensive solution.

Correct expectation: not revolutionary change but gradual, steady improvement in democratic accountability. The timeline is years, not months. The impact is cumulative, not instantaneous.

10.5 Long-Term Governance Vision

Teisond is designed to be built by a founder and ultimately governed by citizens. This is not an aspirational footnote – it is a structural commitment that shapes architectural decisions from the first line of code. Civic accountability infrastructure that permanently depends on a single company, a single founder, or a single jurisdiction contradicts its own premise: that power should be subject to the oversight of those it affects.

The transition from founder-operated to citizen-governed proceeds through three horizons, each building on the stability achieved by the previous one.

The first horizon is operational proof. The Platform launches, activates countries, accumulates data, demonstrates methodological integrity, and achieves financial sustainability through subscription revenue. During this phase, AGPT Ltd operates as sole decision-maker on all matters – methodology, pricing, publication rules, partner relationships, country activation. This concentration of authority is necessary for execution speed at the startup stage, but it is understood as temporary. The governance commitments made in the Neutrality Charter and the Sovereignty and Trust

Framework are the first constraints on founder authority – public, binding, and enforceable from day one.

The second horizon is participatory governance. As the Platform matures and its user base grows, governance authority over non-operational decisions progressively extends to citizens. The scope is deliberately narrow at first: publication conventions, transparency note formats, labelling standards, community guidelines – decisions that affect how the Platform communicates with users, not how it processes data or calculates indices. Citizen participation in these decisions is structured through a governance mechanism that gives verified Platform users the ability to propose, deliberate, and decide on matters within the defined scope. The technical form of this mechanism – whether implemented through distributed ledger technology, token-based governance, or a simpler deliberative structure — is determined by regulatory clarity, technical maturity, and community readiness at the time of implementation, not by a predetermined technology choice.

What does not enter citizen governance scope, at any stage: personal data handling, index calculation methodology, operational security, pricing structure, and compliance decisions. These remain under professional management – first by the founder, later by whatever institutional structure succeeds the founder – because they require technical expertise, regulatory accountability, and the ability to act swiftly in response to security incidents or legal obligations.

The third horizon is institutional independence – described in full in Section 9 (Civic Ownership Architecture). The Platform transitions from company-operated infrastructure to a protocol owned and maintained by its users: distributed data storage on citizen devices, tokenised governance rights, and a community-driven pathway from utility to security token conversion. The destination is fixed – citizen-governed civic infrastructure – and the constraints that apply throughout the journey are inviolable regardless of governance structure: aggregates-only publication, privacy by construction, and methodological neutrality as permanent architectural properties.

The founder's role in this trajectory is specific: build the Platform to the point of operational stability and institutional credibility, design the ownership architecture to accommodate progressive decentralisation, and then step back. The measure of success is not whether the founder retains control, but whether the infrastructure is resilient enough to function – and to be trusted – without any single person at its centre.

10.6 Why Centralised Multi-Tenant Architecture

The architectural choice to operate all 27 EU deployments as a single multi-tenant platform – rather than a federation of independent national platforms – is not a technical convenience. It is a governance decision with direct implications for the Platform's credibility and mission integrity.

Data sovereignty is preserved through localised hosting: each country's citizen data is stored within its own jurisdiction or the EEA, ensuring GDPR compliance without requiring separate legal entities. Consistent methodology across all deployments is guaranteed by a single codebase – there is no possibility of national operators deviating from published methodology, because there are no national operators. Rapid cross-country scaling is enabled by config-driven deployment: launching a new country requires a configuration file, not a development project.

AGPT Ltd as sole data controller simplifies regulatory relationships, eliminates coordination overhead between independent entities, and provides a single point of accountability. If something goes wrong in any deployment, there is one entity responsible, one methodology to audit, and one set of standards to enforce. This clarity is an asset for regulators, researchers, and citizens alike.

10.7 Business Model Aligned with Mission

The Platform's financial sustainability derives from an obvious and psychologically natural source: officials monitoring themselves. The same people whom citizens judge want to see their own legitimacy indices, track dynamics, and compare themselves with colleagues. This creates direct correspondence between business model and mission: the very act of monitoring public legitimacy generates value officials are willing to pay for.

This fundamentally differs from extractive models where user interests conflict with commercial objectives. In Teisond's model, every participant benefits from the same data: citizens gain a voice, officials gain feedback, media gain a data source, researchers gain an empirical layer – and the Platform sustains itself from subscription revenue generated by the data citizens freely provide. No advertising, no data sales, no grant dependence, no conflict of interest between what the Platform publishes and how it earns money.

The pricing structure reinforces this alignment: fixed fees by authority level, country-adjusted for accessibility, with zero relationship between payment and index outcome. A subscribing official and a non-subscribing official with identical citizen judgments display identical public indices.

10.8 Governance as Continuous Practice

Governance principles – independence, neutrality, transparency, privacy protection, accessibility – are operational commitments embedded in technical architecture and contractual relationships, not declarations (see Section 6).

Privacy by construction ensures individual judgments remain private through technical impossibility of profiling, not merely policy prohibition. The Platform inevitably faces pressure to sacrifice independence for finances, tweak methodology for political advantage, or weaken privacy for analytical convenience. Principles provide guidelines; sustaining them requires judgment, resilience, and institutional culture.

The Platform applies to itself the same visibility discipline it applies to officials: methodology changes are documented with public changelogs, financial summaries are published annually, third-party audits will be conducted as the Platform scales, with results disclosed publicly (see § 7.7). An organisation that demands accountability from others must practise it internally.

10.9 Realistic Timeline and Expectations

Creating civic infrastructure requires patience. The Platform launches with progressive country activation in 2026 as countries meet the three convergence criteria described in § 7.2. By Year 2, all activated Platforms are operational and approaching financial self-sustainability. By Year 5, the goal is full coverage across 27 EU member states, with legitimacy monitoring recognised as a legitimate accountability mechanism alongside elections, courts, media, and civil society. By Year 10, expansion beyond the EU – to 40–60 democracies covering one to two billion people – positions the Platform as standard democratic infrastructure. Beyond Year 10, the generational effect begins:

officials who grew up under continuous legitimacy monitoring govern differently from the generation that knew only episodic accountability.

The pre-launch and initial activation phases are self-funded by the founder. Scaling beyond initial operations – security certification, EU entity establishment, intellectual property protection, and operational runway – requires angel-stage external capital, as detailed in § 7.8. The automation-first operational model and sublinear cost scaling mean that financial sustainability is achievable at modest initial adoption rates – the Platform does not require mass adoption to survive its first years.

10.10 Acknowledged Risks

Citizen participation may initially be insufficient for meaningful indices, positioning the Platform as an interesting pilot rather than an authoritative source. Official subscription rates may differ from projections, requiring additional financing or revised commercial terms. Manipulation attempts may episodically distort indices, undermining trust and requiring enhanced safeguards. Political opposition may complicate launch in certain jurisdictions, restraining network potential. Privacy breaches – however unlikely given architectural protections – could seriously undermine user trust, requiring transparent response. Alternative initiatives with better funding may fill the niche before Teisond establishes itself. Mission drift under financial or political pressure remains a constant temptation.

These risks are inherent in any attempt to create new civic-tech infrastructure. They cannot be eliminated – only managed through careful design, honest disclosure, continuous adaptation, and mission discipline. The correct attitude is not confidence that risks will not materialise, but preparedness for when they do.

10.11 Call to Action

For Investors and Strategic Partners: This is civic infrastructure at the formation stage – a platform designed for the entire EU market with a clear revenue model, mission-aligned business logic, and a Blue Ocean positioning in Public Legitimacy Analytics. The risk-return profile combines social impact with commercial viability. Read Sections 4, 5, and § 7.8 to assess the opportunity.

For Officials: This is not a threat – it is a career management tool. Legitimacy indices give you what no other instrument provides: continuous, verified feedback from the citizens you serve. Early subscribers gain insight before public indices become widely cited. Read § 2.3 and Appendix C to understand how the Platform works and what protections you have.

For Media: Go beyond the "ratings + scandals" paradigm. Replace assumptions with verified data. Tell political stories through the language of legitimacy. Make NOLI and office+period scorecards front-page metrics – a new system for public analytics.

For Researchers and Academia: Enter a new discipline at its formation stage. Make Public Legitimacy Analytics a living laboratory where theories of accountability and trust are tested on data. Set the academic standard for methodology in this field.

For NGOs and Civil Society: Engage as partner, user, and advocate. Support sustainable infrastructure of civic judgment – instead of investing in episodic bursts of media, petition, or street emotions. In the digital age, this is more effective, more reliable, and safer for participants.

For the Sceptical: Read on. This document is designed to withstand scrutiny, not to avoid it.

10.12 Final Reflections: Democracy as Continuous Practice

Elections are primary democratic infrastructure. Courts protecting rights, free media investigating abuses, civil society conducting change – all basic infrastructure. The civic institution of governmental legitimacy monitoring aspires to join this list, supplementing existing elements with a component that closes accountability deficits they do not cover.

The vision: a society in which citizens can declare their attitude toward any official whose authority is exercised on their behalf; where judgments of the governed aggregate into public indices which the governing ignore at their peril; where accountability operates continuously; where the question "does this official retain citizens' trust?" receives a clear, verified answer rather than remaining speculation until the next elections.

The accountability deficit exists in all democracies. The question is whether societies will accept it as an inevitable feature or begin building infrastructure for its elimination. Teisond is less a single product than a missing layer of democratic infrastructure: a mechanism for making continuous consent of the governed as visible in the twenty-first century as elections made consent visible in the nineteenth.

This White Paper transitions from documentation to implementation. The concept is articulated. The methodology is specified. The architecture is designed. The legal structure is established. The governance framework is committed. What remains is execution – and the participation of citizens, officials, media, researchers, and partners who share the conviction that democratic accountability should not be confined to election day.

APPENDIX A: THEORETICAL AND PHILOSOPHICAL FOUNDATIONS

A.1 Introduction: Why Theory Matters

This Appendix lays out the theoretical footing of Teisond: what legitimacy means as a measurable construct, why its continuous measurement is both possible and necessary, and how these ideas shape product design, governance choices, and publication rules.

Theory is not ornament. Without clear concepts, "legitimacy" collapses into popularity or partisanship. Theory prevents that collapse by distinguishing legitimacy from approval or single-issue satisfaction, clarifying who is judged (officials exercising governmental authority), and specifying what can be measured (citizen judgments) and how results should appear (aggregated indices with thresholds and confidence intervals).

The Platform's architecture follows from these first principles. Aggregates-only publication protects persons while revealing patterns – no individual records exposed, indices published only above minimum sample thresholds, every published value accompanied by a confidence interval (§5.4). Universality of scope means any office wielding governmental authority belongs in scope, not just high-profile positions. Continuity in time means measurement is ongoing – monthly by default, weekly where volumes allow – because legitimacy drifts between elections and, for appointed officials, no electoral test exists at all. Neutrality by construction means guardrails prohibit political profiling, micro-segmentation below thresholds, and outputs enabling targeted manipulation.

Classical accounts of legitimacy centre on systems and offices, but citizens interact with officials. A citizen's willingness to accept decisions depends on judgments about a specific official's trustworthiness, fairness, and competence. Measuring this personal legitimacy safely and in aggregate connects lived experience with democratic theory.

Clear concepts also discipline governance choices. When anomalies appear – bursts, coordination signals – the Platform discloses flags rather than silently erasing data, because public reason produces better outcomes than hidden curation. If legitimacy belongs to citizens, they should eventually help steward the standards by which it is measured – but only for non-operational governance matters, never for access to personal data (§6.6.1). And because norms travel but laws vary, GDPR serves as the baseline localised upward, with roles and responsibilities mapped explicitly (§6.2).

The appendix proceeds as follows: §A.2 establishes the conceptual framework for legitimacy (Weber, Easton, personal legitimacy as continuous variable); §A.3 identifies two fundamental lacunae in electoral legitimacy (distortion and remission); §A.4 reconceptualises social control as horizontal civic power; §A.5 examines the crisis of civic cohesion; §A.6 synthesises how Teisond addresses these theoretical gaps; §A.7 outlines the civic ownership philosophy behind potential DAO and tokenisation; and §A.8 concludes with how theory informs practice. Throughout, the operative principle is that theory is the contract the Platform makes with readers about what its numbers mean – and what they will never be used to do.

A.2 Legitimacy: Conceptual Framework

Legitimacy, as this White Paper uses the term, is the publicly recognised right to rule – the belief that an official's exercise of authority is appropriate, justified, and should be complied with. It is not a synonym for popularity or policy agreement; it concerns citizens' acceptance of authority over time.

This definition draws on a long tradition in political theory. The sections below trace three foundational contributions – Weber's typology, Easton's distinction between system and authority support, and the treatment of personal legitimacy as a continuous variable – and show how each informs Teisond's design decisions.

A.2.1 Max Weber's Typology

Weber identified three classic sources of legitimacy:

- Traditional legitimacy – acceptance grounded in custom, continuity, and inherited roles.
- Charismatic legitimacy – devotion to a leader's perceived extraordinary qualities; intense but fragile and episodic.
- Legal-rational legitimacy – obedience to impersonal rules and offices; authority attaches to the office, not the person.

Teisond operates within a legal-rational order: the Platform measures citizens' recognition of the authority of officials, not the constitutional validity of the system itself. Results are presented at the office level (personal to the incumbent, yet framed by the office), published only in aggregated form and with statistical guardrails.

A.2.2 David Easton's Three Types of Legitimacy

For democratic systems, it is useful to distinguish:

- Ideological (principled) legitimacy – belief in the regime's core values and purposes.
- Structural (procedural/institutional) legitimacy – confidence in the "rules of the game" (elections, courts, administration).
- Personal (authorities) legitimacy – acceptance of particular officials based on perceived trustworthiness, fairness, and competence.

Teisond's primary object is personal legitimacy of officials (authorities). The Platform remains neutral on ideological disputes and partisan outcomes. This separation keeps measurement focused and prevents category errors: a citizen who distrusts a specific mayor is not necessarily expressing dissatisfaction with democracy itself. Easton's framework makes this distinction analytically precise.

A.2.3 Personal Legitimacy as a Continuous Variable

Elections are episodic; legitimacy drifts between them – and for appointed officials, no electoral test exists at all. Teisond treats personal legitimacy as a time-varying quantity that can strengthen or erode month by month in response to performance, conduct, and events.

The Platform measures this through a deliberately simple mechanism. Verified citizens provide binary signals – personal trust or distrust toward a specific official. These signals are one-citizen-one-account and aggregated per office and period into a scaled index (0–100) with confidence intervals. Publication thresholds protect privacy and quality through minimum sample sizes and rounding. Sub-threshold entities display "Not enough judgments" rather than a number.

The resulting index represents the current public acceptance of authority for a specific official. It does not represent general job performance, policy wisdom, or legal validity, and it is not a profiling tool or a vehicle for micro-targeting. Guardrails enforce this distinction by design: aggregates-only publication with no individual-level exposure, no per-user histories, and no demographic correlations of opinions; anomaly detection with transparency, where suspected bursts or coordination are flagged rather than silently suppressed; and neutral APIs that serve only aggregated outputs, rejecting micro-segments below thresholds.

A continuous, office-level measure of personal legitimacy, presented with statistical and privacy guardrails, gives democratic systems an instrument for visible, non-violent accountability that operates continuously across every level of governmental authority – without crossing into profiling or partisanship.

A.3 Electoral Legitimacy: Two Fundamental Lacunae

Elections are essential but episodic. Between elections, two gaps undermine the informational quality of democratic control: distortion of political preferences at the ballot box and remission of accountability afterward. These are not failures of democratic design but structural limitations inherent in episodic mechanisms. Teisond addresses both by providing a continuous, privacy-preserving signal of personal legitimacy of officials.

A.3.1 Distortion of Political Preferences

Ballots bundle issues, identities, and future contingencies into a single episodic choice. This produces structural distortions. Voters must accept a package – party, coalition, persona – in which specific judgments about particular officials get lost. Strategic voting skews sincere preferences through fear of "wasting" a vote, spoiler dynamics, and tactical considerations. Campaigns compress information into short windows, so that performance signals appearing after election day cannot be expressed until the next cycle. And media salience and polarisation effects can overwhelm fine-grained judgments of trustworthiness and fairness for individual officials.

The consequence is that the electoral result is a coarse, bundled proxy for legitimacy. It tells the public who holds office, not how citizens currently accept or contest the authority of each official. Teisond separates the signal: verified citizens can continuously express trust or distrust toward specific officials, with results aggregated at the office-and-period level (0–100 index with confidence intervals) and published only above privacy and quality thresholds.

A.3.2 Remission of Accountability

After inauguration, accountability tends to remit – it thins out and becomes indirect. Formal checks such as audits, committees, and courts are slow and often opaque to citizens. Officials control the

visibility, timing, and framing of decisions between elections. And protests, petitions, or investigative journalism, while important, are episodic and costly – they do not yield a routine, quantitative signal.

The result is that authority can drift away from public acceptance long before the next election – and for appointed officials, no election ever arrives to correct the drift. No shared metric makes this erosion visible. Teisond restores a low-friction, non-violent accountability channel: a continuous, privacy-preserving legitimacy index per office that can strengthen or erode month by month, with anomaly flags and public transparency notes rather than silent suppression.

A.3.3 Why This Matters for Democratic Theory

Continuous legitimacy measurement matters for democratic theory on several levels. It re-centres citizens: legitimacy is a relationship between citizens and officials, and a continuous, office-level measure makes that relationship visible without exposing individuals. It bridges representation and responsibility: elections authorise, but ongoing acceptance sustains that authorisation, and visibility of the drift disciplines conduct. It improves learning and trust: early signals enable course corrections by officials and more informed public discourse, reducing the cost of errors and polarisation. And it sets principled guardrails: because legitimacy is public rather than personal data, the Platform enforces aggregates-only outputs, minimum sample thresholds, rounding, and no profiling or micro-segmentation below thresholds.

Electoral signals decide who governs; they do not continuously indicate how that authority is being accepted – and for appointed officials, they indicate nothing at all. Teisond supplies the missing instrument: a privacy-safe, statistically disciplined measure of personal legitimacy that operates continuously across every level of governmental authority.

A.4 Social Control Reconceptualised

"Social control" is often heard as coercion or censorship. This appendix reclaims the term in its original civic meaning: the horizontal capacity of citizens to make public power responsive through non-violent visibility and shared standards – without surveillance, profiling, or manipulation. The term has accumulated negative connotations (§A.4.1), but its original meaning (§A.4.2) describes precisely what Teisond provides: a structured, non-coercive mechanism through which citizens collectively influence the informational environment in which authority operates.

A.4.1 Negative Connotations and Historical Baggage

The term "social control" carries baggage that must be acknowledged before it can be reclaimed. It is not coercion – not police power, moral policing, or content takedowns; the Platform exercises no authority over officials but provides information, and what officials do with that information is their decision. It is not stigma or shaming – Teisond never publishes individual citizen data, and official indices are statistical aggregates, not personal verdicts. And it is not partisan mobilisation – outputs are neutral aggregates produced by uniform methodology applied to all officials regardless of affiliation.

A.4.2 Original Meaning: Horizontal Civic Power

In its original sociological sense, social control refers to the mechanisms through which communities maintain coherence and hold power accountable – not from above (coercion) but from alongside (horizontal feedback). Teisond operationalises this meaning.

Citizens continuously signal acceptance or withdrawal of acceptance – trust or distrust – toward specific officials. The signal is private at the individual level and public only in aggregate. Making these signals visible disciplines officials through reputational incentives rather than force: an official whose legitimacy index declines faces questions from peers, media, and constituents without any formal sanction being imposed. And shared, published rules – thresholds, confidence intervals, anomaly flags – turn diffuse sentiment into legible, auditable public information. This is the difference between "people are unhappy" (vague) and "this official's legitimacy index declined from 64 to 51 over three months with 95% CI" (precise, verifiable, actionable).

A.4.3 The Democratic Paradox

Democracies promise rule by the people yet often lack routine, low-friction instruments for the people to exercise continuously. This paradox is structural, not incidental. Institutional checks – courts, audit bodies, ombudsmen – exist but are slow, specialised, and often opaque to citizens; they serve institutional accountability, not citizen-facing accountability. Street politics – protests, petitions – matter but are episodic, high-cost, and available only to the motivated and organised; the ordinary citizen with a grievance about a local official has no practical channel. The result is that citizens' day-to-day leverage is weak, and authority can drift from acceptance with no shared metric making the drift visible.

Teisond resolves the paradox by providing a continuous, privacy-preserving channel where citizens shape the informational environment of governance without targeting individuals. It is not the only solution to the paradox – but it is the specific instrument missing from the current democratic toolkit.

A.4.4 Teisond as Social Control Infrastructure

The Platform operationalises social control as horizontal civic infrastructure through several interlocking design properties. The object of measurement is the official occupying a specific office; the publication unit is the office+period (title and jurisdiction), with signals aggregated into a 0–100 legitimacy index accompanied by confidence intervals. Privacy is enforced by construction: aggregates-only publication, minimum sample thresholds, rounding, no per-user histories, no demographic breakdowns, and no micro-segments below thresholds. Integrity operates through disclosure rather than suppression – anti-manipulation systems detect bursts and coordination, and when they do, users see public flags and transparency notes rather than silently altered data. APIs serve only aggregated outputs; attempts to query below thresholds are rejected regardless of the caller's authentication level.

The civic outcomes follow from this design. Officials gain timely, non-defensive feedback. Citizens and watchdogs gain a shared reference point. Media gain a standard, auditable metric that reduces rumour and spin. Reclaimed as horizontal civic power, social control means visibility rather than force – and the Platform makes that visibility safe, neutral, and statistically disciplined, so that authority remains answerable to the public continuously, across every level of governmental authority.

A.5 Crisis of Civic Cohesion

Democratic systems increasingly lack a shared informational spine: citizens inhabit fragmented media spaces; trust erodes; and routine, low-friction accountability between elections is weak. This fragmentation is not merely a media problem – it is a structural challenge to the democratic premise that citizens share a common factual basis for holding authority accountable. Teisond responds by creating a neutral, privacy-preserving signal of the personal legitimacy of officials that different publics can rely on without exposure or profiling.

A.5.1 Systemic Fragmentation

The informational environment in which democratic accountability operates has fractured along several dimensions. Social feeds and niche media produce non-overlapping realities – the same event appears as different "facts" depending on which information ecosystem a citizen inhabits. Speed and virality outrun careful reasoning, so that rumour and spin dominate the time window in which opinions form; by the time corrections arrive, the narrative is set. And coarse partisan frames crowd out office-level judgments of trustworthiness and fairness – citizens are asked to pick sides, not to judge officials.

The result is that public debate lacks a common, auditable reference for how citizens are currently accepting or withdrawing acceptance of specific officials. Everyone has opinions; no one has shared, verified data.

A.5.2 Decline of Horizontal Power

While institutional checks on authority have grown more sophisticated over time, citizens' direct horizontal leverage has arguably weakened. Day-to-day tools are episodic – protests, petitions, exposés – and high-cost, requiring organisation, publicity, and sustained effort that most citizens cannot sustain. Formal checks are slow, specialised, and often invisible to the public: a citizen who knows that an ombudsman investigated a complaint gains nothing if the outcome is sealed. And without shared, periodic measures of public acceptance, officials' reputations are shaped by media narratives and partisan framing rather than by structured citizen input – authority can drift away from acceptance with no mechanism making the drift visible.

The result is that civic cohesion frays as frustration accumulates without a legitimate, non-violent outlet. Citizens feel powerless; officials feel unaccountable; the distance between them widens.

A.5.3 Why Teisond Matters in This Context

In this fragmented landscape, the Platform matters because it provides a common reference point. A neutral, office-and-period legitimacy index (0–100 with confidence intervals) anchors discourse across media, watchdogs, and civil society – different publics can read the same number the same way. Continuity replaces episodic shocks: monthly (or weekly, where volumes allow) updates produce visible trajectories rather than isolated data points.

Safety is embedded by design: aggregates-only outputs, minimum sample thresholds, rounding, no per-user histories or demographic correlations, and no micro-segments below thresholds. When suspected manipulation is detected, the Platform responds with public flags and plain-language

notes rather than silent takedowns. And incentive alignment ensures that visibility rewards responsiveness and corrective action by officials, while citizens gain a low-friction channel to register acceptance or withdrawal of acceptance without exposure.

Cohesion requires shared, trusted signals. The Platform supplies one for the acceptance of authority: a privacy-preserving, statistically disciplined, office-level measure that different publics can read the same way – reducing rumour, lowering temperature, and restoring continuous accountability.

A.6 Synthesis: How Teisond Addresses Theoretical Gaps

The preceding sections identified three theoretical gaps: electoral distortion (§ A.3.1), accountability remission (§ A.3.2), and the erosion of horizontal civic power (§ A.4–A.5). Teisond supplies the missing, privacy-preserving instrument for continuous accountability of officials. It resolves the two electoral gaps and operationalises social control as horizontal civic visibility, not coercion. This section synthesises how the Platform's design maps to each gap.

A.6.1 Continuous Measurement of Personal Legitimacy

The core mechanism is simple in concept and disciplined in execution. Verified citizens express trust or distrust toward specific officials. These signals are aggregated per office and period into a 0–100 index with confidence intervals. Indices appear only above minimum sample thresholds; sub-threshold entries display "Not enough judgments." Rounding protects anonymity near thresholds, and anomaly flags ensure disclosure over suppression.

The outcome is the first continuous, verified, office-level measure of public acceptance of authority available in democratic practice – a timely, comparable legitimacy signal that operates across every level of governmental authority, not only between elections but for officials who never face elections at all.

A.6.2 Closing the Remission Gap

Citizens do not wait years to express acceptance or withdrawal of acceptance – the Platform provides low-friction feedback that operates continuously. Officials see month-by-month movement and can address causes early, enabling course correction rather than discovering erosion only at the next election – or, for appointed officials, never discovering it at all. And a shared, auditable metric reduces rumour, selective framing, and spin by grounding public discourse in verifiable data.

Schedler's view of accountability rests on two pillars: answerability (public reason-giving) and enforcement (consequences for abuse). Without a continuous public signal, accountability thins out between electoral cycles – and for the vast majority of officials who never face elections, it never thickens at all. The Platform supplies the missing informational layer – neutral, office+period indices read by citizens, civil society, media, and oversight bodies – while leaving enforcement to existing legal and institutional channels.

A.6.3 Enabling Horizontal Social Control (Non-Violent, Neutral)

The Platform enables horizontal social control by maintaining a clear separation between measurement and publication. The official is the object of measurement; results are published at the office+period level (title and jurisdiction), with no personal records exposed. Privacy is enforced by construction: aggregates-only outputs, no per-user histories, no demographic correlations, and no micro-segments below thresholds. When manipulation is suspected, the Platform flags anomalies and publishes plain-language notes rather than silently erasing signals. APIs expose only aggregated endpoints; disallowed breakdowns are rejected.

Legitimacy is a scarce coordination resource: communities allocate power and attention to actors and rules they perceive as legitimate. To sustain that legitimacy, the Platform prioritises process over ad-hoc fixes – clear publication rules, aggregates-only outputs, transparency notes, and a Right to Respond. Participation is encouraged without profiling; any future governance evolution remains standard-focused and excludes personal data.

A.6.4 Addressing Democratic Fragmentation

The Platform addresses democratic fragmentation by providing a common reference point: a neutral, statistics-first index that anchors debate across media, watchdogs, and civil society. Participation is distributed and continuous – many citizens contributing judgments toward individual offices, not siloed by platform tribes. And incentive alignment ensures that visibility rewards responsiveness and penalises neglect without targeting individuals.

These properties rest on design commitments documented elsewhere in this White Paper: data governance including thresholds, anonymity, and publication policy (§5.4); controller and processor mapping with explicit guardrails (§6.2); and a governance pathway ensuring that any future citizen co-governance remains standard-focused and excludes personal data (§6.6.1).

The boundaries are equally important. The Platform measures public acceptance of authority for each official over time, safely and comparably. It does not profile individuals, enable targeted mobilisation, adjudicate legal validity, or replace elections or courts. By turning citizens' trust and distrust into a continuous, privacy-safe, office-level signal, the Platform restores continuous accountability without force or exposure – filling the structural absence that elections alone cannot address and that, for appointed officials, no existing mechanism addresses at all.

A.7 DAO and Tokenisation: Civic Ownership Philosophy

Purpose. This section articulates the theoretical foundation for civic ownership of legitimacy monitoring infrastructure. The question it addresses is: if legitimacy monitoring infrastructure serves the public, should the public have a role in governing and owning it? The answer is yes – within carefully scoped guardrails that protect privacy, neutrality, and methodological integrity. Section 9 translates these theoretical commitments into a concrete civic ownership architecture, describing the mechanisms through which the Platform transitions from founder-operated infrastructure to a protocol owned and maintained by its users.

A.7.1 From Representation to Participation to Ownership

Democratic accountability evolves through three stages, each building on the previous. Representation through elections authorises office-holders but leaves long intervals with weak citizen leverage – and for appointed officials, provides no leverage at all. This is the starting point: necessary but insufficient. Participation through the Platform restores a continuous, privacy-preserving legitimacy signal about specific office-holders, giving citizens a voice that operates permanently rather than episodically. Ownership goes one step further: citizens and public-interest actors help steward the standards by which the Platform operates – labels, thresholds, transparency cadence – within strict guardrails. Ownership here means stewardship of public rules and artefacts, not equity or revenue. It is the principle that infrastructure serving citizens should ultimately be governed by citizens.

A.7.2 Civic Data as Commons (Without Exposing Persons)

Teisond produces a specific kind of public good: aggregated indices, confidence intervals, threshold notices, and transparency notes. These outputs are the commons – shared, reusable, and valuable to media, researchers, civil society, and citizens alike.

The commons excludes personal data categorically. There are no individual records in public or subscriber outputs; no per-user histories; no demographic correlations; no micro-segments below thresholds. Publication discipline – minimum sample sizes, rounding, and anomaly flags – ensures that the commons serves the public interest without compromising any individual (see § 5.4).

This distinction is crucial: the Platform produces public goods from private inputs, and the architecture guarantees that the transformation is irreversible. There is no path from published aggregates back to individual judgments.

A.7.3 Tokenisation as Alignment Mechanism (Governance-Utility Only)

The governance-utility token (U-token) is conceived as an alignment mechanism, not a financial instrument. It functions as a key for scoped participation – enabling verified users to propose, discuss, and decide on non-personal-data matters such as transparency note cadence and non-material parameter bounds, all within the guardrails defined in §5.4. The U-token is not a security, not a subscription, and not equity: it carries no dividends, no promise of returns, no access to personal data, and does not substitute for paid subscriptions.

Rollout is compliance-first. No public token sale is planned at this stage; pilots may use non-transferable allocations in compliant jurisdictions, and decisions or hashes can be anchored on-chain while execution remains off-chain (§6.6.1). Anti-capture safeguards include quorums, proposal cooldowns, rate-limits, optional delegation with transparent records, and key revocation for abuse.

A.7.4 Decentralised Governance and Democratic Values

Decentralised governance, if and when it emerges, must align with the Platform's democratic values. Neutrality means governance concerns method and transparency, not political outcomes or profiling. Subsidiarity means national participation comes first, with any network-level layer remaining advisory for shared standards such as taxonomies and open components. Auditability means material decisions carry human-readable changelogs, and where on-chain hashes are used,

they provide public audit trails. And accountability symmetry means the same visibility discipline the Platform applies to officials applies to itself – rules change only with notice, reasons, and records.

Civic ownership for Teisond means stewardship of shared standards and public artefacts, enforced by privacy and neutrality guardrails. A governance-utility token and carefully scoped participatory processes can align incentives and widen participation – without exposing individuals, promising financial returns, or departing from the aggregates-only design (§§5.4 and 6.6.1).

A.8 Conclusion: Theory Informs Practice

Democracies need a continuous, privacy-preserving way to see whether citizens accept the authority of specific office-holders – not only between elections, but permanently, including for the vast majority of officials who never face elections at all. Teisond provides that instrument without profiling, force, or partisanship.

The mechanism is straightforward. Verified citizens register trust or distrust toward named office-holders. Signals are aggregated per office and period into a 0–100 legitimacy index with confidence intervals, published only when sample thresholds are met; otherwise readers see "Not enough judgments."

This matters for three reasons. It resolves two structural deficits in electoral accountability: distortion at the ballot box (bundling, strategy, timing) and remission of accountability afterward – while extending coverage to officials who fall outside the electoral mechanism entirely. It reclaims social control as horizontal civic visibility: reputational discipline through transparent, aggregated information rather than force. And it gives citizens, civil society, watchdogs, journalists, and public institutions a shared reference point that lowers rumour and spin.

Guardrails are embedded by design. Publication is aggregates-only: no individual records, no per-user histories, no demographic correlations, no micro-segments below thresholds. Rounding protects anonymity near boundaries, and suspected coordination is flagged with public notes rather than silently suppressed. APIs expose only aggregated endpoints; disallowed breakdowns are rejected. Data governance and responsibilities are mapped explicitly (§§5.4 and 6.2), and any future citizen co-governance remains standard-focused and excludes personal data (§6.6.1).

The index should be read for what it is: a statistically disciplined signal of current public acceptance of authority for an office-holder. It is not a prediction of elections, a verdict on policy wisdom, or a legal ruling – and it should always be read together with its confidence interval and publication notes. This White Paper covers legitimacy of authority only; extensions to additional civic signals remain out of scope until they can meet the same privacy, neutrality, and statistical standards.

Theory in this appendix is not ornament. It is the contract behind the product. On that basis, Teisond turns civic judgment into a safe, neutral, and continuous public signal – so that authority remains answerable to the governed at every level and at all times.

APPENDIX B: GLOSSARY OF TERMS

Terms used throughout the White Paper. Where helpful, entries cross-reference §§5.4 (Data Governance and Compliance), 6.2 (Roles and Jurisdictions Passport), and 6.6.1 (Potential DAO Integration).

Aggregates-only – a design principle: the Platform publishes only aggregated statistics (never individual records), with privacy and quality guardrails (thresholds, rounding). See §5.4.

Alert (threshold crossing / significant move) – a notification that an index crossed a configured threshold or moved by a material amount over a period. Alerts are aggregated events; no personal data.

Anomaly flag – a public marker shown when volumetric/behavioral anomalies are detected (e.g., bursts, suspected coordination). Favors disclosure over suppression.

API (Aggregated-only API) – public/subscriber endpoints that return only aggregated outputs (indices, benchmarks, alerts). Disallows micro-segments below thresholds; rate-limited and auditable. See § 5.4.10.

Below-threshold / Sub-threshold – a state where sample size n is below the minimum required for publication. The UI shows "Not enough judgments" instead of a number.

Benchmark – an aggregated comparison across a set of offices/regions within a period (e.g., top-N lists). Sub-threshold entities are excluded.

Citizen verification (privacy-preserving) – mechanisms that enforce one-citizen-one-account semantics without collecting personal identifiers. Implementation specifics are omitted to prevent gaming; the WP specifies guarantees and architectural properties.

Confidence interval (CI) – an uncertainty band around the index (e.g., 95% CI) reflecting sample size and variance. Every published index includes a CI.

Controller / Processor / Sub-processor – GDPR roles: AGPT Ltd is Controller for national platform data; cloud infrastructure providers are Sub-processors under appropriate data processing agreements. See § 6.2.3.

Data governance – the set of policies and technical measures enforcing privacy, quality, and lawful processing: aggregates-only publication, thresholds, access controls, logs minimisation, and retention rules. See § 5.4.

DPA (Data Processing Agreement) – contractual terms governing Processor obligations, TOMs, breach notices, deletion/return, and audits.

Judgment – a single trust/distrust signal submitted by a verified citizen regarding a specific office-holder within a period. Judgments are aggregated; no per-user histories are stored.

AGPT Ltd – Advanced Global Polling Technology Ltd (UK). The company operating all national Teisond deployments. Acts as data Controller for each national platform. Responsible for legal

compliance, platform operations, and publication policy enforcement across all jurisdictions. See § 6.2.1.

Governance-utility token (U-token) – a non-financial access key for scoped participation in standards/transparency processes (create/propose/decide within guardrails). Not equity, not a subscription, not a promise of returns. See § 6.6.1 and A.8.

Index (Legitimacy Index) – a 0–100 measure of current public acceptance of authority for a specific official, published at the office+period level with a confidence interval.

k-anonymity guardrails – practical safeguards to ensure that published aggregates cannot be used to infer information about small groups or individuals (via thresholds and rounding).

Managed Hosting – the operational model in which AGPT Ltd directly operates all hosting infrastructure for every national deployment through a centralised multi-tenant architecture. AGPT Ltd acts as data Controller. See §§ 5.4 and 6.2.

Micro-segmentation (disallowed) – any breakdown or query that would fall below publication thresholds or enable re-identification. Rejected by the API and UI.

"Not enough judgments" – the standardised UI message for below-threshold cases; prevents false precision and protects privacy.

n (sample size) – the count of judgments aggregated for an office+period. Publication requires meeting minimum sample thresholds set per country (privacy & quality). See § 5.4.

National DAO (future, scoped) – a token-gated process for non-PII governance matters (labels, transparency cadence, non-material parameter bounds). Advisory and compliance-bounded. See § 6.6.1.

Network DAO (future, advisory) – a cross-country standards layer (taxonomies, transparency formats, open components). No PII; advisory unless ratified contractually. See § 6.6.1.

Office – publication unit (office+period: title and jurisdiction). The object of measurement is the official occupying the office; results are published at the office+period level with privacy guardrails (thresholds, rounding). See: Official; Office-level (publication unit); Index; §5.4.

Official - The person currently exercising governmental authority in a specific office; the object of measurement for the legitimacy index. Publication remains at the office+period level; no individual records are exposed.

Office-level (publication unit) – the level at which Teisond publishes results: aggregated metrics for a specific office+period (e.g., "Mayor of Riverton – September 2025"). Not an individual record; subject to thresholds and rounding.

One-citizen-one-account – a verification property ensuring each natural person can contribute at most one account, enforced without collecting personal identifiers.

Period – the temporal window for aggregation (default monthly; weekly where volumes allow). Each period produces a fresh index per office.

Privacy by construction – architectural stance that prevents collection/storage of personal identifiers and disallows profiling by design; privacy is not an add-on.

Publication thresholds – minimum conditions (e.g., n ≥ 100) required before an index is shown publicly; includes rounding.

Public (free) vs Subscriber (paid) – public outputs: headline aggregated indices, CIs, threshold notices, transparency notes. Subscriber outputs: aggregated benchmarks, historical slices, alert feeds, and aggregated API access. No row-level data in either.

Rounding – publishing indices at one decimal place (or coarser, if configured) to reduce re-identification risk and false precision.

Roles & Jurisdictions Passport – a living artifact per country that maps roles (controller/processor/sub-processor), lawful bases, local constraints, and escalation trees. See § 6.2.

SLA (Service Level Agreement) – availability, latency, incident response targets, and remedies for platform operations. Defined internally by AGPT Ltd as part of operational standards.

Sub-processor – a processor engaged by the Processor to carry out specific processing activities (e.g., cloud/infra), bound by sub-DPA terms and no secondary use.

Subscriber – a paying organisation (e.g., media, NGO, research, administrations) with access to aggregated dashboards, alerts, and API, within guardrails.

Threshold notice – a visible label that communicates insufficient sample size instead of publishing a number.

Transparency note – a plain-language public explanation accompanying material changes, anomalies, or method updates; part of the "disclosure over suppression" norm.

Trust / Distrust (input semantics) – the only MVP signals citizens give about officials. These are personal judgments of acceptance or withdrawal of acceptance of authority, aggregated at the office-and-period level.

APPENDIX C: FREQUENTLY ASKED QUESTIONS

Practical answers for citizens, officials, media, researchers, and regulators. See also Appendix B (Glossary), Section 2 (Concept and Methodology), Section 5 (Legal Structure), and Section 6 (Governance and Ethics).

C.1 General

Q1. What does Teisond actually measure?

A. Verified citizens register trust or distrust toward specific officials. These civic judgments are aggregated and published as a 0–100 legitimacy index for each office+period (e.g., "Mayor of Riverton – September 2025"), with a confidence interval.

Q2. Is this a popularity poll or a rating?

A. No. The index reflects public acceptance of authority (legitimacy), not job performance, policy approval, or personal likability. It is a structured civic judgment, not an opinion survey.

Q3. Do you predict election results?

A. No. The index is not a polling average or a model of electoral intention. It should always be read with its confidence interval.

Q4. Who benefits from this signal?

A. Citizens gain a continuous channel for accountability. Officials gain timely feedback on public acceptance. Media, civil society, watchdogs, and researchers gain a shared, auditable reference point that operates continuously.

Q5. Who operates Teisond? Is it independent?

A. Teisond is operated by AGPT Ltd, a UK-registered technology company. AGPT Ltd has no political affiliations, no governmental funding, and no partisan investors. The Platform's revenue comes exclusively from official subscription fees (see C.5). The Neutrality Charter, published on every national platform, commits Teisond to algorithmic transparency and non-interference.

Q6. Who finances the Platform?

A. Development is self-funded by AGPT Ltd. Ongoing operations are financed through B2G/B2B subscriptions – officials and institutions pay for enhanced analytics. Citizens use the Platform for free. There is no advertising, no data sales, and no political sponsorship.

Q7. In which countries is Teisond available?

A. Teisond launches simultaneously across all 27 EU member states, each with a dedicated national platform ({country}.teisond.com) in the local language. Full functionality is activated in waves, determined by which countries first meet activation criteria: verification infrastructure connected, official database populated, and sufficient waitlist demand. Priority markets include Estonia, Netherlands, Poland, Spain, and Germany. Other countries have landing pages with waitlist registration from day one.

C.2 For Citizens

Q8. How do I participate?

A. Register on your national platform ({country}.teisond.com) and verify your identity through the Platform's verification process (government-issued document check with biometric liveness confirmation). Where national eID systems are connected, they are available as an additional verification option. After verification, you can express trust or distrust toward any official in the system.

Q9. Is participation free?

A. Yes. Registration, identity verification, and expressing civic judgments are entirely free. Paid features are available only for officials and institutional subscribers who want enhanced analytics.

Q10. How is my anonymity guaranteed?

A. Your identity data is processed through one-way cryptographic hashing (SHA-256 with a seasonal salt). The Platform stores only the resulting fingerprint – never your name, ID number, or any personal identifier. Even platform developers cannot determine who you are. The cryptographic salt rotates annually, and after two years the old salt is permanently deleted, making any retroactive identification technically impossible.

Q11. Can I change my judgment?

A. Yes. Your current judgment (trust, distrust, or neutral) toward any official can be updated at any time. However, to prevent manipulation, changes are rate-limited. Only your current judgment counts – the system does not store a history of your changes.

Q12. What happens if I delete my account?

A. Your account and personal fingerprint are permanently removed. Your past judgments remain in the aggregated statistics (they contributed to published indices), but they can never be traced back to you – they are part of anonymous aggregate data only.

Q13. What are my rights under GDPR?

A. You have the right to access your data, the right to erasure (account deletion), and the right to lodge a complaint with your national data protection authority. Because Teisond uses irreversible hashing, most personal data is not stored in identifiable form. For data subject requests, contact privacy@teisond.com.

C.3 Method & Privacy

Q14. Do you measure the office or the person?

A. Object of measurement: the official currently holding the office. Publication unit: the office+period (title and jurisdiction). No individual citizen records are exposed.

Q15. What happens if there are not enough judgments?

A. Below the minimum sample threshold (default n ≥ 100, adjustable per country), the UI shows "Not enough judgments." No number is published until the threshold is met.

Q16. What are the privacy guardrails?

A. Aggregates-only outputs; no per-user histories; no demographic correlations; no micro-segments below thresholds; rounding to reduce false precision. The Platform is designed as privacy-by-construction: personal data is never collected in identifiable form.

Q17. How do you prevent manipulation?

A. Since only document-verified citizens with biometric liveness confirmation can participate, bot attacks are structurally impossible. The real threat is coordinated campaigns by real people. Integrity systems detect bursts, temporal clustering, geographic concentration, and account-age anomalies. The Platform favours disclosure over suppression: suspicious patterns are flagged publicly with a plain-language transparency note. Detection specifics are not published to prevent gaming.

Q18. What is the period of measurement?

A. Monthly by default; weekly where volumes allow. Each period yields a fresh office-level index with a confidence interval.

C.4 Reading and Using the Index

Q19. How should journalists cite the metric?

A. Use: Office, Period, Index, CI (e.g., "Mayor of Riverton, Sep-2025: 63.4 (95% CI 61.9–64.9)"). Avoid over-interpreting single-month noise; discuss trajectory and confidence. Do not use the index as an election forecast.

Q20. What does a decline mean?

A. Erosion of public acceptance of authority. It is a signal for explanation, engagement, or course correction – not a legal verdict or a recall mechanism.

Q21. Can I compare across regions or levels?

A. Yes, via benchmarks (aggregated comparisons). Ensure comparable periods and note differences in sample sizes and confidence intervals.

Q22. Do you publish breakdowns by age, gender, or party?

A. No. To prevent profiling and re-identification, no demographic breakdowns are provided at any access level.

C.5 Product & Access

Q23. What's public vs. paid?

A. Public (free, registered users): headline legitimacy index (%), trend arrow, sample size. Subscriber (paid, officials and institutions): absolute trust/distrust counts, consensus index, time-series with hourly detail, comparative analytics, anomaly reports. No row-level or individual data at any level.

Q24. Is there an API?

A. Yes – aggregates-only. Endpoints expose indices, benchmarks, and alerts with rate-limits and audit trails. Queries below publication thresholds are rejected.

Q25. How are official subscriptions priced?

A. Fixed pricing by authority level (L1–L4), adjusted per country for accessibility. Pricing is structured to avoid any conflict of interest: fees are fixed regardless of the index value, so there is no incentive for the Platform to influence results. The subscription provides access to analytics – it does not affect the index in any way.

C.6 Officials, Media & Civil Society

Q26. I'm an official – how should I use this?

A. Track your trajectory with confidence intervals. Read transparency notes. Address causes of declines through responsiveness and engagement. Legitimacy is sustained by accountability, not messaging alone.

Q27. Can an official contest the index?

A. Yes. Use the Right to Respond channel: your official statement is published as a transparency note for the relevant office+period. Recalculations occur only under the revision policy (material error, confirmed bug, or verified manipulation). Individual citizen data is never disclosed.

Q28. Can an official "opt out" of monitoring?

A. No. Public offices are objects of legitimate public interest. Publications are aggregated with privacy thresholds. Officials retain the Right to Respond and can use subscriber dashboards to understand and address causes of change.

Q29. Can a new office be added to monitoring?

A. Yes. Criteria: the position involves exercise of governmental authority, sufficient expected participation, and no conflict of interest. New offices appear with "Not enough judgments" until the sample threshold is reached.

Q30. I'm a watchdog/NGO – what can I do with it?

A. Use benchmarks and alerts to prioritise oversight, surface early-warning signals, and hold reason-giving dialogues with offices showing sustained declines. Subscriber access provides historical data for longitudinal analysis.

Q31. I'm a journalist – what are common mistakes?

A. Treating the index as an election horse-race; ignoring the confidence interval; cherry-picking single months; inferring motives without evidence; asking for disallowed micro-segments; and presenting the index without noting sample size.

C.7 Compliance & Operations

Q32. Who is the data Controller under GDPR?

A. AGPT Ltd is the data Controller for all national platforms. Cloud infrastructure providers operate as Sub-processors under data processing agreements. See § 6.2.

Q33. Where is citizen data stored?

A. Citizen data for each country is stored in jurisdiction-appropriate infrastructure within the EU, ensuring GDPR data residency compliance. National deployments are isolated tenants – a breach in one country's infrastructure does not compromise others.

Q34. What gets published if there's a suspected manipulation attempt?

A. The aggregated index remains (subject to thresholds), accompanied by an anomaly flag and a plain-language transparency note.

Q35. Do you ever remove or revise published numbers?

A. Material changes follow publication policy: rationale, public notice, and versioned records. Revisions are documented, not hidden. The revision history is part of the public transparency archive.

Q36. How do I exercise my GDPR rights or file a complaint?

A. Contact privacy@teisond.com for data subject requests (access, erasure, portability). You also have the right to lodge a complaint with your national data protection authority. Response timelines comply with GDPR requirements (30 days).

C.8 Governance Evolution

Q37. Will there be decentralised governance or tokens?

A. In the long term, Teisond may introduce a governance-utility mechanism allowing citizens and public-interest actors to participate in decisions about standards, transparency processes, and non-personal-data matters. This is not a financial instrument, not equity, and not a substitute for platform access. No token sale is planned at this stage. See § 6.6.1 for the governance evolution roadmap.

Q38. How do you prevent governance capture?

A. Any future governance participation operates within strict guardrails: quorums, proposal cooldowns, rate-limits, transparent delegation, and revocation for abuse. Governance processes never touch personal data, profiling, or index methodology. Core publication rules remain under AGPT Ltd authority as mission protection.

APPENDIX D: BIBLIOGRAPHY AND REFERENCES

Citation style: APA 7th (concise). Purpose: sources that inform the White Paper's theoretical footing, measurement choices, privacy guardrails, integrity controls, and governance stance. Inclusion ≠ endorsement of any policy position.

D.1 Core Theory of Legitimacy and Democratic Accountability

- Weber, M. (1978). Economy and Society (G. Roth & C. Wittich, Eds.). University of California Press. – Classic typology (traditional/charismatic/legal-rational).
- Easton, D. (1965). A Systems Analysis of Political Life. Wiley. – Regime support vs support for authorities; basis for focusing on officials.
- Easton, D. (1975). A Re-assessment of the Concept of Political Support. British Journal of Political Science, 5(4), 435–457. – Refinement of "support" signals over time.
- Beetham, D. (1991). The Legitimation of Power. Palgrave. – Legitimacy as a normative and empirical concept (used here to separate acceptance from popularity).
- Dahl, R. A. (1971). Polyarchy: Participation and Opposition. Yale University Press. – Accountability and participation between elections.
- Pitkin, H. F. (1967). The Concept of Representation. University of California Press. – Representation vs ongoing responsiveness (motivates continuous signals).
- Schedler, A. (1999). Conceptualizing accountability. In A. Schedler, L. Diamond, & M. F. Plattner (Eds.), The self-restraining state: Power and accountability in new democracies (pp. 13–28). Lynne Rienner. – Classic treatment distinguishing answerability and enforcement; supports the need for a continuous, inter-electoral accountability signal focused on officials.

D.2 Measurement and Statistical Guardrails

- Agresti, A. (2013). Categorical Data Analysis (3rd ed.). Wiley. – Foundations for proportion estimates and uncertainty.
- Agresti, A., & Coull, B. A. (1998). Approximate is better than "exact" for interval estimation of binomial proportions. The American Statistician, 52(2), 119–126. – CI choice for aggregated binary inputs.
- Clopper, C., & Pearson, E. S. (1934). The use of confidence or fiducial limits... Biometrika, 26(4), 404–413. – Reference for conservative intervals.
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 15. – Taxonomy for integrity/flagging systems.

D.3 Privacy by Design and Publication Thresholds

- Sweeney, L. (2002). k-Anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5), 557–570. – Basis for thresholds and anti-reidentification.
- Dwork, C. (2006). Differential Privacy. In ICALP 2006 (pp. 1–12). Springer. – Foundational privacy framework; informs the Platform's approach to threshold-based publication safeguards.
- European Union. (2016). Regulation (EU) 2016/679 (GDPR). – Roles (Controller/Processor), lawful bases, data-minimization; see §§ 5.4 & 6.2.
- NIST. (2020). SP 800-53 Rev. 5: Security and Privacy Controls. – Catalog for security controls (hosting/SLA alignment).
- ISO/IEC 27001. (2022). Information Security Management. – Governance benchmark for managed hosting.

D.4 Human Rights and Public Interest Basis

- United Nations. (1948). Universal Declaration of Human Rights (Art. 19). – Expression and access to information.
- UN General Assembly. (1966). International Covenant on Civil and Political Rights (Arts. 19, 25). – Participation in public affairs; expression.
- Council of Europe. (1950). European Convention on Human Rights (Art. 10). – Freedom of expression.

D.5 Neutrality, Transparency and Media Use

- Tufte, E. R. (2001). The Visual Display of Quantitative Information (2nd ed.). Graphics Press. – Clarity over spin; avoid deceptive visuals.
- Groves, R. M., et al. (2009). Survey Methodology (2nd ed.). Wiley. – Limits of episodic polling (motivates continuous, low-friction signals).
- Meyer, P. (2002). Precision Journalism (4th ed.). Rowman & Littlefield. – Standards for reporting quantitative findings (CI, uncertainty, reproducibility).
- Council of Europe, Committee of Ministers. (2004). Recommendation Rec(2004)16 on the right of reply in the new media environment. – Normative basis for a Right to Respond and transparent notes, favoring disclosure over takedowns while protecting individuals.

D.6 Governance, Tokens and DAO (Non-Financial, Scoped)

- Buterin, V. (2021). The Most Important Scarce Resource is Legitimacy. (essay). – Legitimacy in decentralized governance (conceptual backdrop; not financial).

- Hassan, S., & De Filippi, P. (2021). Decentralized Autonomous Organizations. Internet Policy Review, 10(2). – Overview of DAO governance patterns and risks.
- Wright, A., & De Filippi, P. (2015). Decentralized Blockchain Technology and the Rise of Lex Cryptographia. – Early articulation of on-chain governance boundaries.

Note: Token here is governance-utility only (non-financial, no PII), per § 6.6.1 and Appendix A.8.

D.7 How to Cite this White Paper

- In text: "Teisond White Paper (vX.Y, 2025-MM-DD), Section N.M."
- Bibliography entry: Teisond White Paper, vX.Y (2025-MM-DD). AGPT Ltd. CC BY-ND 4.0.

D.8 Versioning Note

This bibliography is maintained with the method. When publication rules, thresholds, or security controls materially change, the bibliography and cross-references are updated in the next minor version.

APPENDIX E: GOVERNMENTAL AUTHORITY – DETAILED POSITION ESTIMATES BY LEVEL

This Appendix provides detailed breakdowns of official positions at each of the four authority levels described in Section 2.2. Estimates are based on publicly available institutional data from EU member states and represent ranges reflecting the diversity of administrative structures across the EU-27. These figures serve as reference material for database construction and coverage planning; the Platform's actual database is populated progressively from official registries and institutional directories.

All estimates refer to a typical mid-sized EU member state unless otherwise noted. Actual counts vary significantly by country size, administrative tradition, and institutional structure.

E.1 Level 1: National Authority

Executive Branch (60–80 officials per country)

Head of State and/or Head of Government: 1–2 positions. Deputy Heads of Government: 1–4 positions, varying by system. Cabinet Ministers: 15–25 ministers leading major government departments, including Interior/Home Affairs, Defence, Economy/Finance, Foreign Affairs, Justice, Education, Health, Transportation, Environment, Labour, and Agriculture. Junior Ministers and Secretaries of State: 30–50 senior political appointees; each ministry typically has 1–3 junior ministers handling major policy areas. Heads of major national agencies (politically appointed): 10–20 officials, including intelligence services, transportation authorities, national railways, airports authority, postal services, and broadcasting corporations.

Legislative Branch (520–560 officials per country)

Lower chamber (primary legislative body): 300–400 members in a mid-sized country. Upper chamber (if bicameral): 200–250 members. Legislative leadership – speaker, deputy speakers, committee chairs: 30–50 officials. Party leaders and parliamentary group leaders: 8–15 officials.

Other National Positions (5–15 officials per country)

Supreme or Constitutional Court justices (where appointments are political): 10–15 justices. Ombudsman and similar oversight institutions: 1–3 officials. Heads of major regulatory agencies (if politically appointed): 3–8 officials, including securities regulator, competition authority, telecommunications regulator, and energy regulator.

Level 1 Summary

Typical count per EU country: 250–800 officials, depending on population and institutional structure. EU-27 total: approximately 12,000–15,000 national-level officials. National official counts scale sublinearly with population – legislative chambers do not triple when population triples: Estonia's Riigikogu has 101 members; Germany's Bundestag has 736.

E.2 Level 2: Regional Authority

Regional Executives (50–70 officials per country)

Regional heads (governors, regional presidents, premiers): 12–18 officials. Deputy regional heads and senior councillors: 30–50 officials. In federal and devolved systems, these include regional equivalents of cabinet ministers; in centralised systems, senior administrators.

Regional Legislators (1,200–1,300 officials per country)

Regional assembly or state legislature members: 1,100–1,200 representatives. Size varies – larger regions may have 100+ members, smaller regions 50–80. Regional assembly leadership and committee chairs: 50–100 officials.

Provincial/Intermediate Governance (where applicable, 50–100 officials per country)

Some member states maintain an intermediate tier between regional and municipal levels: provincial governors, council leaders, or provincial administrative directors.

Level 2 Summary

Typical count per EU country: 400–2,500 officials, depending on administrative structure. Federal systems (Germany, Austria, Belgium) have substantially more regional officials per capita than centralised unitary systems. The number of first-tier administrative subdivisions ranges from 3 (Belgium) to 16 (Germany's Bundesländer) to 21 (France's metropolitan regions).

E.3 Level 3: Municipal Authority

Municipal Structure Variation

Municipal structure shows the greatest variation across the EU. Some countries have many small municipalities (Spain with 8,132; France with 35,000+); others have consolidated larger units (Denmark with 98 after municipal mergers). Municipal autonomy also varies: Nordic countries grant strong municipal authority; centralised systems retain more control at regional and national levels.

For a mid-sized EU country, the approximate distribution is: major cities (over 500,000 population): 5–10; mid-size cities (100,000–500,000): 30–60; small cities (20,000–100,000): 200–400; towns (5,000–20,000): 1,000–2,000; small towns (under 5,000): 2,000–3,000.

Municipal Executives

Mayors: one per municipality, totalling 3,000–5,000 officials. Deputy mayors and municipal managers: approximately 1,000–2,000 officials.

Municipal Legislators (Councillors)

Total varies enormously based on municipal structure and council size regulations. Major cities: 25–60 councillors each (150–600 total). Mid-size cities: 15–30 each (450–1,800 total). Small cities: 11–21 each (2,200–8,400 total). Towns: 7–15 each (7,000–30,000 total). Small towns: 5–9 each (10,000–27,000 total). Potential total: 20,000–70,000 municipal legislators in a mid-sized country.

Municipal Administrators

Municipal secretaries and managers: 2,000–4,000. Department directors in larger municipalities: 1,000–3,000.

Level 3 Summary

Typical count per mid-sized EU country: 25,000–80,000 officials (wide range depending on municipal structure). EU-27 total: approximately 800,000–2,000,000 officials.

E.4 Level 4: Local Officials and Appointed Administrators

Education Sector (15,000–25,000 officials per country)

School principals and directors (public schools): 10,000–18,000 officials across primary and secondary schools. Each principal exercises authority over disciplinary decisions, teacher assignments, curriculum implementation, school policies, budget allocation, and parent-teacher relationships. School district administrators and inspectors: 500–1,500 officials. University and college administrators: 200–500 officials (where publicly managed). Regional education directors: 200–500 officials.

Law Enforcement and Public Safety (8,000–15,000 officials per country)

Local and municipal police chiefs: 3,000–6,000 officials, exercising authority over traffic enforcement, community policing, response protocols, and personnel deployment. National or state police local commanders: 1,000–3,000 officials. Fire department chiefs: 1,000–2,000 officials. Emergency services directors: 500–1,500 officials. Prison directors (where publicly managed): 200–500 officials.

Regulatory and Inspection Services (10,000–20,000 officials per country)

Building inspectors and urban planning officials: 3,000–7,000 officials, with authority over permit approvals, building code compliance, and construction authorisation. Health and safety inspectors: 2,000–4,000 officials. Environmental inspectors: 1,000–3,000 officials. Labour inspectors: 1,000–2,000 officials. Consumer protection officials: 500–1,500 officials. Licensing and permit officials: 2,000–4,000 officials.

Social Services and Public Welfare (8,000–15,000 officials per country)

Social services directors and supervisors: 4,000–8,000 officials, with authority over benefit eligibility, service allocation, case management, and emergency assistance. Child welfare

supervisors: 1,000–3,000 officials. Public housing administrators: 1,000–2,500 officials. Disability services administrators: 1,000–2,000 officials. Employment service directors: 1,000–2,000 officials.

Healthcare Administration (3,000–7,000 officials per country)

Public hospital directors: 500–1,200 officials (where healthcare is publicly managed). Health centre directors: 2,000–5,000 officials. Regional health service directors: 50–150 officials. Emergency medical services directors: 200–800 officials.

Administrative and Financial Services (8,000–15,000 officials per country)

Tax office local directors: 1,000–3,000 officials. Social security and pension office directors: 1,000–3,000 officials. Vehicle registration office directors: 500–1,500 officials. Land registry officials: 2,000–4,000 officials. Municipal finance directors: 1,000–2,000 officials. Pension and benefits administrators: 2,000–3,000 officials.

Judicial and Legal Services (5,000–10,000 officials per country)

Local and regional court judges: 4,000–8,000 judges across civil, criminal, administrative, and family courts, with authority over case decisions, sentencing, and custody determinations. Magistrates and justices of the peace: 1,000–3,000 officials. Public prosecutors at local level: 500–1,500 officials.

Transportation and Infrastructure (3,000–7,000 officials per country)

Public transit authority directors: 500–1,500 officials. Municipal public works directors: 2,000–4,000 officials. Port authority directors: 50–150 officials. Airport directors: 30–100 officials (where publicly managed).

Environmental and Utilities Services (2,000–5,000 officials per country)

Water utility directors: 500–1,500 officials. Waste management directors: 500–1,500 officials. Parks and recreation directors: 500–1,500 officials. Environmental protection officers: 500–1,500 officials.

Cultural and Community Services (3,000–6,000 officials per country)

Library directors: 1,000–2,000 officials. Museum directors: 500–1,200 officials. Community centre directors: 1,000–2,000 officials. Sports facility directors: 500–1,500 officials.

Level 4 Summary

Typical count per mid-sized EU country: 50,000–150,000 officials. EU-27 total: approximately 1,500,000–3,000,000 individuals – by far the largest category of officials exercising governmental authority.