Funded by
the European Union

Supporting regional entrepreneurship through the adoption of innovative technologies, including AI, in public services

Drawing on lessons learned from
Andalusia, Extremadura, Madrid, Navarra (Spain)
and North Rhine-Westphalia (Germany)

eGA
e-governance academy

Funded by
the European Union

Project
EU Commission Project (24ES06/24DE33)

Prepared for
Directorate-General for Structural Reform Support (now-SG REFORM)

Implemented by
e-Governance Academy

Beneficiary regions:
North Rhine-Westphalia (Germany)
Andalusia (Spain)
Extremandura (Spain)
Madrid (Spain)
Navarra (Spain)

Authors
Federico Plantera, Piret Hirv

Contributors
Andres Vimb, Kalle Tomingas, Konstantyn Krupskyi, Diletta Huyskes, Mykhailo Kornieiev, Viljar Peep, Radu Serrano

This document was produced with the financial assistance of the European Union via the Technical Support Instrument and implemented by the e-Governance Academy, in cooperation with the European Commission.

The views expressed herein can in no way be taken to reflect the official opinion of the European Union.

eGA
e-governance academy

Published by the e-Governance Academy
Ahtri 6, 10151, Tallinn, Estonia
ega.ee

Contents

List of abbreviations

Abbreviation	Full term
AI	Artificial intelligence
API	Application programming interface
EU	European Union
FAST	Fairness, accountability, sustainability, transparency
GDPR	General Data Protection Regulation
IPA	Intelligent process automation
IT	Information technology
NLP	Natural language processing
OECD	Organisation for Economic Co-operation and Development
RAG	Retrieval-augmented generation
SMEs	Small and medium-sized enterprises
SWOT	Strengths, weaknesses, opportunities, threats
XAI	Explainable AI

Executive summary

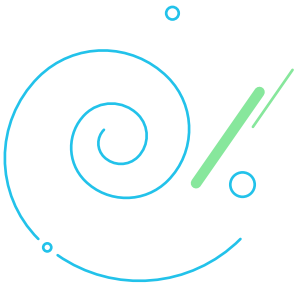

The modernisation of public administration is frequently mischaracterised as a purely technical challenge – a matter of procuring faster servers, sleeker interfaces or more sophisticated algorithms. This perspective, while seductive in its simplicity, obscures the profound institutional and political dimensions of digital transformation. A rigorous analysis of the entrepreneurial ecosystems in project beneficiary regions in Germany and Spain reveals that the primary barrier to innovation is not a scarcity of technology but a legacy of institutional design.

The current landscape is defined by fragmentation: data is trapped in administrative silos, processes are disjointed, and the user experience for the entrepreneur is one of friction (e.g. repetitive paperwork, bureaucracy and absence of one-stop shops) rather than support.

A rigorous analysis of the entrepreneurial ecosystems reveals that the primary barrier to innovation is not a scarcity of technology but a legacy of institutional design.

Within the scope of the present EU project, this report argues for the introduction of innovative technologies – specifically artificial intelligence (AI) – as a fundamental redesign of public administration itself, and not merely as an upgrade to existing IT systems. The strategic pivot is from a model of "digitisation", in which paper processes are simply replicated on screens, to a model of "intelligent service", in which the state proactively anticipates needs, automates compliance and empowers the user. These innovative technologies can enable disruptive service models – fundamentally new ways for entrepreneurs to interact with government – rather than merely making existing processes electronic.

This shift requires technology to be viewed through a socio-economic lens: code is not neutral infrastructure but a form of policy and governance, particularly in contexts where algorithms determine how public services are delivered. The architecture that is built today will define the relationship between citizens and the state for years to come.

Drawing on a comprehensive synthesis of regional expert analyses commissioned ad hoc, legal assessments and global best practices, this report identifies three critical shifts required to achieve this vision:

This report argues for the introduction of innovative technologies – specifically artificial intelligence (AI) – as a fundamental redesign of public administration itself, and not merely as an upgrade to existing IT systems.

1) **From silos to ecosystems**
The current fragmentation of registries prevents the state from seeing the entrepreneur as a whole entity. The priority must be the creation of an interoperability layer that allows data to flow securely between agencies, enabling the implementation of the Once-Only Principle. As shown by the Abundance Institute's analysis[1] and many others, automated data exchange is more than an efficiency measure; it is a prerequisite for trust and economic velocity.

2) **From reactive to proactive governance**
Current services wait for the entrepreneur to apply. AI enables a shift towards proactivity, in which the state uses predictive analytics to offer support (e.g. grant eligibility) before the user even asks for it. This reframes the state from a passive administrator to an active partner in economic value creation, a concept reinforced by the OECD's emphasis[2] on anticipatory innovation.

3) **From compliance to agency**
The ultimate goal of technology is to empower the human user. Whether it is the entrepreneur in rural Extremadura who needs voice-activated guidance or companies in Madrid that seek API access, the technology must lower the barriers to economic participation. Technology must be measured by its capacity to care for user well-being and foster civic trust, as articulated in the Alan Turing Institute's ethics guidelines.[3]

The recommendations set out in this report are grounded in the specific legal, operational and cultural realities of the target regions. They propose a path forward that balances the transformative potential of AI with the strict guardrails required by European values and regulations, so that efficiency does not come at the cost of rights.

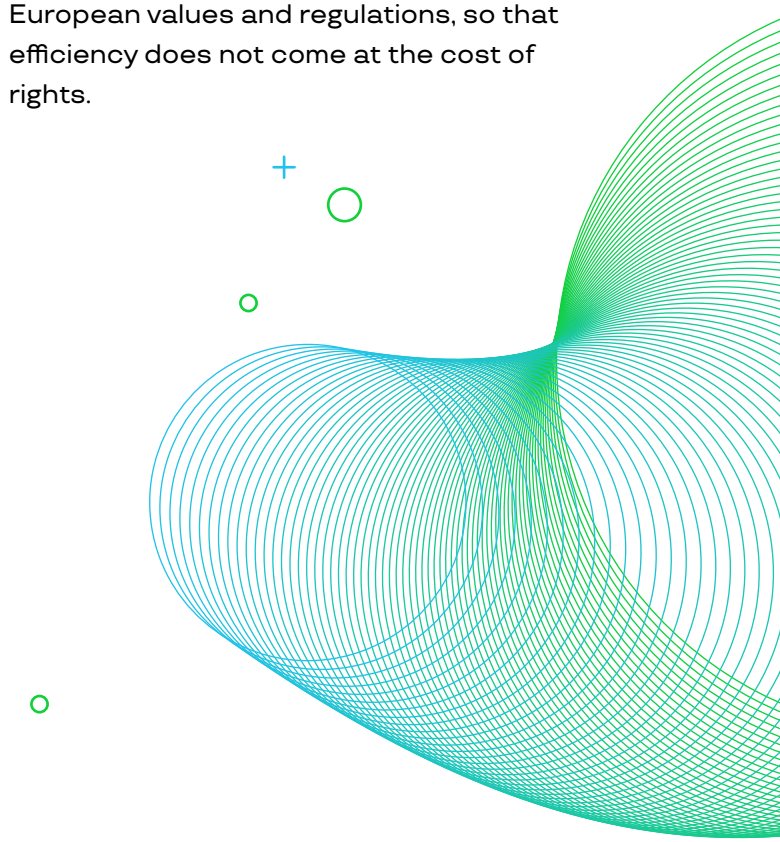

1 Ilves et al. (2025). *Government Digitization Efficiency and Performance in the US and Estonia*. Abundance Institute
2 OECD (2025). *Governing with Artificial Intelligence*.
3 Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector*. The Alan Turing Institute.

Methodology and sources

This report synthesises findings from multiple analytical streams conducted throughout the project. The primary evidence base consists of five regional SWOT analyses that systematically assessed the current state of digital services for entrepreneurs in Andalusia, Extremadura, Madrid, Navarra (Spain) and North Rhine-Westphalia (Germany). These analyses were conducted through structured workshops with beneficiary authorities and supplemented by desktop research into recommendations for service portfolios, institutional needs and technical infrastructure.

In parallel, the project commissioned **seven thematic expert evaluations** covering regional service delivery models, AI solution mapping, digital enablers and interoperability requirements, legal and regulatory frameworks, ethics and risk assessment, cybersecurity considerations, and scalability analysis. These expert inputs were produced by specialists within the e-Governance Academy consortium and provide the technical depth underpinning the technology assessments in Chapters 2 and 3. Where the report references "expert analysis" or "regional analysis", it draws on these commissioned evaluations.

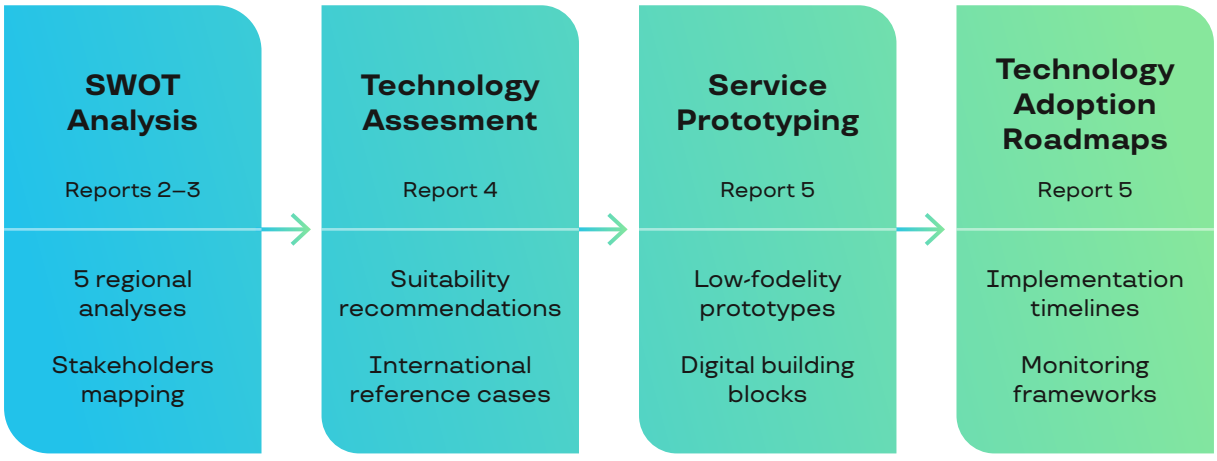

Figure 1. Project methodology flow – from regional analysis to implementation planning

The international reference cases in Chapter 4 were identified through a structured review of digital government implementations, prioritising cases with documented outcomes and transferable lessons for the target regions. External sources – including reports from the OECD,[4] the Alan Turing Institute[5] and the Abundance Institute[6] – provide broader contextual evidence and are cited throughout.

4 OECD (2025). *Governing with Artificial Intelligence.*
5 Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.* The Alan Turing Institute.
6 Ilves et al. (2025). *Government Digitization Efficiency and Performance in the US and Estonia.* Abundance Institute

Key concepts

This report uses several technical terms that are essential to understanding the recommendations. For readers less familiar with digital governance terminology, the following definitions provide practical context.

Interoperability refers to the ability of various IT systems, databases and agencies to exchange and use information seamlessly. In practice, this means that when an entrepreneur registers a business with one agency, that information automatically becomes available to tax authorities, social security and licensing bodies – without the entrepreneur having to re-submit the same documents multiple times. Achieving interoperability typically requires agreed data standards, secure exchange protocols and institutional agreements that define who can access which data under which conditions.

Prerequisites (or digital building blocks) are the foundational infrastructure elements that must be in place before advanced technologies like AI can function effectively. These include reliable and clean data systems, digital identity verification, basic cybersecurity protections and clear data governance rules that define who owns and manages different datasets. Without these foundations, deploying AI is like installing a sophisticated navigation system in a car with no engine – the technology exists but cannot deliver value.

Guardrails are the rules, safeguards and oversight mechanisms that ensure AI systems operate fairly, transparently and within legal boundaries. In practice, guardrails include requirements such as human review before automated decisions become final, audit trails that document how a recommendation was generated, mechanisms that allow citizens to understand and challenge algorithmic outcomes, and regular testing to detect bias or errors. Guardrails exist both to maintain public trust and to comply with regulations such as the EU AI Act.

The Once-Only Principle is an EU policy commitment to require citizens and businesses to provide the same information to public authorities only once. If a company registration document has already been submitted to one agency, it should not have to be submitted again to another. Implementing this principle requires interoperability between systems – the technical and institutional arrangements that allow agencies to share data securely rather than asking users to act as couriers of their own information.

Decision support describes AI systems that provide recommendations, analysis or options to a human official, who then makes the final decision. This contrasts with **decision-making** systems in which the AI determines outcomes directly (e.g. automatically approving or rejecting a

grant application). This distinction has significant regulatory consequences: under the EU AI Act, systems that autonomously determine eligibility for public benefits are classified as high-risk and trigger substantial compliance obligations. For most regions, remaining in decision-support mode is the prudent path.

High-risk AI is a classification under the EU AI Act (Annex III) that applies to AI systems used in sensitive contexts, including systems that determine eligibility for public services and benefits. High-risk classification triggers mandatory requirements, such as conformity assessments, human oversight protocols, detailed documentation and ongoing monitoring. Systems that remain advisory (*decision support*) rather than determinative (*decision-making*) can often avoid this classification.

Legacy systems are the existing IT systems, databases and software that agencies currently use. They are often built years or decades ago on older technology. These systems typically still function for their original purpose but are difficult to modify, integrate with other systems or extend with new capabilities. Most regions face the challenge of deploying new AI tools that must work alongside, rather than replace, these legacy systems.

API (application programming interface) is a standardised way for different software systems to connect and communicate with each other. When an API-first approach is recommended, this entails designing systems so that data and services can be accessed programmatically by other systems – not only through a human-facing website. APIs are the building blocks that enable interoperability: they allow a business portal to pull data from a tax registry, or a chatbot to check licensing requirements, without requiring manual data transfer.

Chapter 1:

The landscape of digital services for entrepreneurs

1.1 The institutional reality: Fragmentation as statecraft

To understand the potential for AI in public services, one must first surface the substantial reality of the administrative state as it exists today. Across the regions analysed, the primary barrier to innovation is not a lack of technology, but a specific architectural legacy: fragmentation. This is the digital reflection of analogue statecraft, in which authority has historically been distributed across distinct, autonomous silos.

A Deloitte analysis[7] of the same regional ecosystems in Spain found "excessive complexity and fragmentation of the service ecosystem" across all four Spanish regions, with stakeholders reporting that "services and programmes for entrepreneurs are dispersed and lack coordination between agents and resources". The study documented how different administrations operate with "different particular interests that are poorly aligned", creating what interviewees described as a lack of cohesion, whereby the incentives facing the administration diverge from those facing the entrepreneur. The consequence is duplication, gaps and a user experience that forces entrepreneurs to navigate a maze of overlapping programmes without a unified point of entry.

Across the regional analyses commissioned from the project's pool of experts, data trapped in "isolated silos" reflect institutional rigidity and a historical preference for departmental autonomy over systemic coherence. Each agency operates as a fenced-off area, hoarding its data rather than treating it as a common public resource. Systems are frequently characterised as "disjointed", and this disjointedness creates a user experience that forces entrepreneurs to bridge the gaps between government departments. The burden of integration is shifted onto the citizen, who must navigate disconnected portals, re-submit the same information and interact with agencies that operate as if they were entirely separate entities. This fragmented user experience creates significant friction and functions as a *direct tax* on economic productivity.

7 Deloitte (2023). *Technical Analysis Report: Overcoming Barriers for Innovation in the Entrepreneurship Ecosystems of Andalusia, Navarre, Extremadura, and Madrid.* EU Technical Support Instrument / DG REFORM.

The Once-Only Principle – the foundational concept that a citizen should supply data only once – remains technically impossible in many jurisdictions.

Reliance on manually maintained registry systems suggests an approach whereby administrative authority remains tied to specific, static data structures rather than dynamic data flows. Legacy systems, while functional, often struggle to adapt to new business models or cross-sectoral innovation. The consequence, as highlighted in the expert analysis on digital transformation enablers (see Methodology on page 7), is that the Once-Only Principle – the foundational concept that a citizen should supply data only once – remains technically impossible in many jurisdictions. Without a functioning API ecosystem or a data broker layer, the state cannot "see" itself, let alone automate its functions. It operates with a fragmented memory, unable to recall what it already knows about its citizens. This mirrors a Nortal whitepaper's observation[8] that "highly complicated, fragmented processes and systems are harder to change or improve", which increases the risk of failure and stifling innovation.

The contrast with global best practices is stark. An Abundance Institute comparison[9] of the United States and Estonia shows that transaction costs are dramatically reduced when the state is designed as a platform in which 99% of public services are online and information exchange is automated, such as through the X-Road data exchange layer. In a single year, X-Road saved more than 2,589 years of working time.[10] In Estonia, a company can be established in minutes because the underlying registries are interoperable by design.

This is not only about speed; it is about trust. A unified system signals competence and respect for the citizens' time. In the target regions for this study, the absence of such interoperability transforms what should be a simple digital transaction into a complex bureaucratic journey, eroding trust and creating friction that stifles entrepreneurial activity. The OECD reinforces this point,[11] noting that "interoperability via data exchange platforms and cross-border integration" (or cross-agency integration, in this case) is a key enabler of efficient digital governance, allowing for the seamless flow of information that underpins modern service delivery.

8 Raieste, A., Solvak, M., Velsberg, O., McBride, K. (2025). *Government Efficiency in the Age of AI: Toward Resilient and Efficient Digital Democracies*

9 Ilves et al. (2025). *Government Digitization Efficiency and Performance in the US and Estonia*. Abundance Institute

10 X-Road Factsheet: https://e-estonia.com/wp-content/uploads/factsheet_x-road.pdf

11 OECD (2025). *Governing with Artificial Intelligence*.

1.2 Mapping capability constraints

If fragmentation is the structural barrier, the scarcity of specialised human capital is the operational constraint. Regional analyses (see Methodology) point to a critical shortage of in-house technical talent, creating a dependency on external vendors that threatens digital sovereignty. This is a governance risk. When public administrations lack the internal capacity to design or even procure complex systems, they risk hollowing out the state's ability to govern its own infrastructure. They become consumers of technology rather than architects of public value.

When public administrations lack the internal capacity to design or even procure complex systems, they risk hollowing out the state's ability to govern its own infrastructure.

This capability gap creates a vicious cycle. Because the administration lacks the skills required to manage complex digital projects, it outsources them to vendors who build proprietary, black-box systems. These systems, in turn, may be difficult to integrate or update, which reinforces the fragmentation described above. Expert inputs highlight this risk, noting that "weak communication" between agencies is often exacerbated by incompatible technical standards imposed by different vendors.

The Alan Turing Institute identifies[12] this skills gap as a primary barrier to AI adoption, noting that without internal expertise, governments struggle to "assess, simplify and redefine" their needs. This leads to suboptimal procurement outcomes in which the technology fails to meet actual user needs.

Moreover, this skills gap is compounded by cultural resistance to change. One of the expert analyses for this project (see Methodology on page 7) notes a strong adherence to existing administrative structures, where the introduction of AI is viewed not as an enabler but as a disruption to established workflows. This aligns with findings in the OECD's *Governing with AI* report,[13] which identifies "organisational and cultural resistance" as a dominant theme in the public sector's slow adaptation to algorithmic tools. The fear is often rooted in a misunderstanding of AI's role, viewing it as a replacement for human judgment rather than an augmentation of administrative capacity. To counter this, the OECD emphasises[14] the need for upskilling and reskilling programmes, noting that "automation bias" and "algorithmic aversion" can be mitigated through targeted training that empowers civil servants to work *alongside* AI, transforming them from data processors into decision validators.

Furthermore, the technical prerequisites for advanced AI – specifically the availability of high-quality, structured data – are often missing. Readiness assessments indicate that while there are pockets of

12 Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.* The Alan Turing Institute.
13 OECD (2025). *Governing with Artificial Intelligence.*
14 ibid.

excellence, no authoritative source for entrepreneurial data exists. As the expert analysis on AI safety commissioned for this project (see Methodology on page 7) warns, "inadequate or skewed data" is a primary vector for algorithmic failure. Training predictive models on fragmented or low-quality data produces noise rather than intelligence. Without a robust data governance framework that ensures accuracy, completeness and representativeness, any attempt to layer AI on top of existing systems will likely amplify existing inefficiencies and biases. The OECD stresses[15] that "robust, quality data is an underlying prerequisite for enjoying AI's benefits", and highlights that many AI projects fail precisely because of data-related issues.

1.3 The human dimension: Entrepreneurial archetypes

To design effective services, it is necessary to look beyond the abstract notion of the "user" and observe the specific user archetypes interacting with public entrepreneurship support services. The expert inputs (see Methodology on page 7) allow for the construction of distinct personas that reveal divergent needs within the ecosystem.

First, there is the **Traditional Artisan** (P02), who often operates in traditional sectors or areas with limited digital infrastructure. For this entrepreneur, the primary barrier is not only bureaucratic complexity but also digital literacy. Before complex dashboards, this group requires trust and clarity. The expert analysis (see Methodology on page 7) suggests that "empathetic guidance" – interfaces that translate administrative jargon into plain language – is essential for this group. If the technology alienates them, it fails, regardless of its sophistication. For the Artisan, technology must be invisible, serving as a bridge rather than a barrier. This aligns with the OECD's finding[16] that "tailored services to address personalised citizen needs" are a key benefit of AI, particularly for users who may struggle with standard digital interfaces.

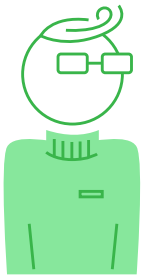

In contrast, the **Digital Explorer** (P01) is highly digitally savvy and typically based in urban hubs, such as Madrid or larger towns in Andalusia and Navarra, and operates with a different set of expectations. This persona demands speed, API access and integration, building on the advanced digital skills that this archetype already possesses. They are frustrated by the siloed nature of current services and seek a sandbox environment in which they can test ideas and access funding without friction. They view the state as a platform provider, expecting the same level of digital seamlessness encountered in the private sector. This reflects the OECD concept of Government as a Platform,[17] in which common building blocks such as digital identity and data exchange enable coherent transformation across the public sector, allowing innovators to build on top of state infrastructure.

15 OECD (2025). *Governing with Artificial Intelligence.*

16 OECD (2025). *Governing with Artificial Intelligence.*
17 ibid.

Finally, the **SME Manager** (P03), whether already established or seeking to establish a company, prioritises efficiency above all and is prevalent in industrial contexts. Their interaction with the state is transactional and compliance-heavy. For this group, the absence of a unified digital identity or predictive compliance tools represents a measurable cost to their business. They require systems that respect their time, offering automated compliance checks and proactive notifications that prevent regulatory friction. This echoes the potential for "continuous compliance" discussed in the OECD report,[18] in which real-time validation replaces periodic reporting, reducing administrative burden and enabling businesses to focus on growth rather than paperwork.

A one-size-fits-all AI strategy will inevitably fail some groups, as the risk of exclusion is not hypothetical.

Understanding these archetypes is critical. A one-size-fits-all AI strategy will inevitably fail some groups, as the risk of exclusion is not hypothetical. A recent UNDP analysis[19] of AI deployment across the Asia-Pacific region warns that AI "does not operate in a homogeneous space but within uneven social, economic, and technological landscapes". Where infrastructure investment concentrates in urban centres, rural and underserved populations become what the report terms "technologically forgotten" – they receive slower response times, lower service quality and less personalisation than their urban counterparts. The study documents how AI-powered health platforms in Indonesia failed to identify rural women as service recipients because of gaps in digital records, how farmers in Bangladesh could not interpret precision agriculture advice delivered through unfamiliar interfaces, and how elderly citizens across multiple countries missed pension payments and welfare applications because system designers failed to account for their needs.

The lesson is clear: without deliberate design for diverse user groups, AI does not close existing divides – it widens them. The task, therefore, is to design systems that are adaptive enough to serve the rural artisans through voice-activated simplicity while offering the urban innovators the high-speed digital infrastructure they require. This requires a nuanced approach to service design that prioritises inclusivity and adaptability over rigid standardisation.

18 ibid.
19 UNDP (2025). AI for the Next Generation of Public Services: *The Next Great Divergence – Why AI May Widen Inequality Between Countries*. Institute for AI International Governance (I-AIIG), Tsinghua University. Available at: https://www.undp.org/sites/g/files/zskgke326/files/2025-12/ai-for-the-next-generation-of-public-services.pdf

Chapter 2:

Overview of innovative technologies

2.1 From deterministic to probabilistic governance

The transition to AI-enabled government marks a shift from *deterministic* automation to *probabilistic* intelligence. This distinction is fundamental to understanding where AI creates value and where it introduces risk. It is a shift from the certainty of rules to the likelihood of patterns.

The strategic error that many governments make is applying the wrong logic to the wrong problem.

Deterministic automation applies rigid logic to structured data. It is the domain of the "if-this-then-that" rules that govern traditional bureaucracy. For processes such as company registration or tax filing, where legality depends on precision, deterministic systems are indispensable. They ensure that rules are applied consistently every time. The Nortal analysis[20] reinforces this view, arguing that rule-based systems should remain the "default" for well-defined administrative tasks in order to preserve transparency and auditability. Attempts to replace these foundational systems with probabilistic models without careful consideration can lead to unpredictability and a loss of trust. The OECD report[21] similarly advises that rule-based automation is best suited to "repetitive and standardised" tasks, ensuring consistency and auditability.

Probabilistic AI, by contrast, deals in patterns and likelihoods. It excels where data is unstructured, rules are ambiguous or the volume of information exceeds human processing capacity. One opportunity identified in some beneficiary regions, namely the use of AI to match entrepreneurs with the most relevant mentors or funding opportunities, is a probabilistic task. There is no single correct mentor, but there is an optimal probability of a good match. Similarly, identifying potential fraud in procurement data is a pattern-recognition task that is well suited to machine learning. The OECD report notes[22] that AI is particularly effective for "detecting data outliers, hidden relationships [...] and other anomalies"

20 Raieste, A., Solvak, M., Velsberg, O., McBride, K. (2025). *Government Efficiency in the Age of AI: Toward Resilient and Efficient Digital Democracies.*

21 OECD (2025). *Governing with Artificial Intelligence.*

22 ibid.

that human reviewers might miss, thereby enhancing accountability.

The strategic error that many governments make is applying the wrong logic to the wrong problem.

Attempting to use probabilistic large language models to interpret complex regulatory requirements, without appropriate guardrails, invites hallucination and legal risk. Conversely, trying to hard-code a recommendation engine for thousands of diverse startups results in a rigid, static listing that quickly becomes outdated. The path forward is a hybrid architecture, with deterministic systems for the binding administrative layers and probabilistic AI for the *advisory* and *navigational* layers. This dual approach allows governments to harness the efficiency of automation while leveraging the insights of AI, ensuring that innovation does not come at the cost of reliability.

2.1.1 Digital enablers: The foundations for AI adoption

Before any AI system can deliver value, certain foundational infrastructure must be in place. These *digital enablers* are prerequisites without which technology investments will fail to produce returns. The regional analyses consistently identify gaps in these enablers as the primary barriers to innovation.

Before any AI system can deliver value, certain foundational infrastructure must be in place.

Interoperability infrastructure: the ability of different agency systems to exchange data seamlessly. Without interoperability, AI tools can only access information within their own silo, which severely limits their utility. Interoperability requires agreed data standards, secure exchange protocols (APIs) and institutional agreements governing data access.

Data governance frameworks: clear rules defining who owns, manages and can access different datasets. This includes data quality standards (ensuring information is accurate, complete and current), retention policies and compliance with GDPR requirements. AI systems trained on poor-quality data will produce poor-quality outputs.

Digital identity systems: mechanisms for authenticating users and authorising access to services. Robust digital identity enables personalised services, secure data sharing with user consent and the implementation of the Once-Only Principle. Weak identity infrastructure forces users to re-authenticate across services and prevents seamless user journeys.

Cybersecurity foundations: the protection of systems and data from unauthorised access, manipulation or disruption. AI systems that process sensitive entrepreneurial data require robust security throughout their lifecycle – from secure development practices to ongoing monitoring for vulnerabilities. Security is not a feature to be added later; it must be designed in from the outset.

Human and institutional capacity: the skills and organisational structures required to procure, implement, govern and evolve AI systems. This includes

technical expertise to manage vendors and evaluate solutions, governance capacity to establish oversight mechanisms and change management capability to support staff through new ways of working. Technology without the human capacity to govern it creates risk rather than value.

The technology recommendations in the following sections must be read against the backdrop of these enablers. Where a region has strong interoperability but weak data governance, certain AI applications become feasible while others remain premature. The Strategic Recommendation Matrix in Chapter 5 maps these enabler profiles to appropriate technology choices.

2.2 Technology suitability assessment

Based on the specific needs identified in the regional analyses, three categories of innovative technology emerge as most suitable for the entrepreneurial ecosystem. However, innovative technologies are dependent on foundational digital building blocks that act as prerequisites:


interoperability

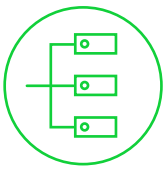
data governance


digital identity


cybersecurity

Where these building blocks are partially developed or fragmented, technology, including AI, use will remain correspondingly constrained in scope and ambition.

As a result, selecting appropriate technologies requires more than matching capabilities to needs. Each technology must be assessed against feasibility within the current digital maturity of beneficiary authorities, compliance with legal and regulatory frameworks such as GDPR and the EU AI Act, organisational readiness and capacity to operate and govern such technologies, and the proportionality of risk in relation to public sector responsibilities. The following sections evaluate each technology category through these lenses.

2.2.1 Conversational AI and virtual assistants

Suitability: high for inclusivity and navigation

The demand for 24/7 information access is universal across the regions, but it is particularly acute where it serves as a bridge across the digital divide. The expert inputs propose conversational AI not merely as a search tool, but as an *empathetic* interface that can translate bureaucratic language into plain terms. This shifts the burden of understanding from the citizen to the system.

- **Application.** A voice-activated assistant could allow rural entrepreneurs to navigate complex grant applications without requiring deep digital literacy. This is a direct response to the needs of the Traditional Artisan persona, lowering the barrier to entry for state

support. *The Agentic State* white paper[23] envisions this as "personalisation and multi-modality", where interactions are customised to language, channel and tone, supporting "universal coverage".

The expert inputs propose conversational AI not merely as a search tool, but as an empathetic interface that can translate bureaucratic language into plain terms.

- **Evidence.** The OECD[24] highlights the success of Portugal's Practical Guide to Justice chatbot, which improved access to information through the use of plain language. Similarly, the Alan Turing Institute notes that, for non-technical stakeholders, the *means of content delivery* is as critical as the content itself. Effective communication is a core component of service delivery. The OECD also points to emerging initiatives such as the UK's Caddy[25] assistant and France's Albert as pilot projects that explore internal tools to support public servants, indicating growing interest in conversational AI across European public administrations.
- **Risk.** The primary risk is *hallucination*, meaning the provision of incorrect information. To mitigate this risk, the system must use retrieval-augmented generation (RAG), grounding its responses strictly in official documents rather than relying on the model's internal training data. This ensures that the guidance provided is accurate and consistent with official requirements. The OECD report[26] and the internal risk assessment conducted for this project (see Methodology on page 7) identify hallucinations as a key operational risk and emphasise the need for careful mitigation to prevent "misguided decisions or actions".

2.2.2 Predictive analytics and decision support

Suitability: high for policy and compliance

For regions with more mature digital governance infrastructure, the expert analysis highlights the need to move from reactive to proactive governance. Predictive analytics offers the ability to anticipate outcomes, whether identifying entrepreneurs at risk of insolvency or modelling the impact of a new subsidy policy. This shifts the state from a passive administrator to an active partner in economic development.

- **Application.** A data-driven governance dashboard could analyse anonymised economic data to forecast the impact of support programmes, allowing for iterative policy adjustments. Predictive models could help SME managers anticipate compliance bottlenecks before they occur, turning compliance from a burden into a manageable process. This aligns with the OECD's observation[27]

23 Ilves, L., Kilian, M., Parazzoli, S. M., Peixoto, T. C., Velsberg, O. (2025). The Agentic State: Rethinking Government for the Era of Agentic AI.
24 OECD (2025). Governing with Artificial Intelligence.
25 Caddy, AI-powered assistant acts as a copilot for customer service agents, empowering them to provide high-quality, actionable advice quickly and securely.
26 ibid.
27 ibid.

For regions with more mature digital governance infrastructure, the expert analysis highlights the need to move from reactive to proactive governance.

that AI can enhance "decision-making and sense-making of the present", allowing for more effective resource allocation and policy targeting.

- **Evidence.** The OECD reports[28] that predictive analytics in the Brazilian courts reduced case analysis time from 44 minutes to just seconds. However, the Alan Turing Institute[29] warns of automation bias, meaning the tendency for humans to uncritically accept algorithmic scores. This highlights the need for training and cultural change alongside technological implementation.
- **Risk.** The primary risk is the black box problem. If a predictive model flags a business as high-risk for a grant, the reasoning must be explainable. As noted in the regional legal analysis (see Methodology on page 7), such profiling triggers GDPR Article 22 protections, which require human intervention and contestability. Transparency is non-negotiable in algorithmic decision-making. The Alan Turing Institute[30] emphasises the importance of "outcome transparency" to ensure that explanations are clear and understandable to non-technical stakeholders.

2.2.3 Intelligent process automation (IPA)

Suitability: high for efficiency

This technology bridges the gap between legacy systems and modern efficiency. It creates a layer of intelligence that sits on top of fragmented databases, automating the manual transfer of data that currently consumes civil service time. It is a pragmatic response to interoperability challenges in the absence of a full system redesign.

- **Application.** Fragmented registry systems require significant manual effort to maintain and update. IPA could automate the classification of businesses based on their descriptions, allowing registries to be maintained dynamically rather than manually. This frees up human resources for higher-value tasks.
- **Evidence.** The Abundance Institute[31] factsheet on Estonia demonstrates that automating back-end processes, such as once-only data exchange, can save up to two per cent of GDP. Efficiency at this scale is transformative. The Alan Turing Institute report[32] supports this view, noting that automating "mundane and repetitive tasks" can significantly boost productivity, enabling civil servants to focus on "mission-critical activities".
- **Risk.** The primary risk is a lack of accountability. If an automated process incorrectly classifies a business,

28 ibid.
29 Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.* The Alan Turing Institute.
30 ibid.
31 Ilves et al. (2025). *Government Digitization Efficiency and Performance in the US and Estonia.* Abundance Institute.
32 Leslie, D. (2019). *Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector.* The Alan Turing Institute.

thereby affecting its eligibility for support programmes, tax obligations or regulatory requirements, the line of responsibility must be clear. The system must be designed in line with accountability-by-design principles, ensuring an audit trail for every automated action. Automation must not lead to an abdication of responsibility. The OECD report[33] identifies "reduced job quality" and "public service workforce displacement" as potential risks if automation is not managed using a human-centric approach.

33 OECD (2025). *Governing with Artificial Intelligence.*

Chapter 3:
The trust framework

3.1 Ethics as part of the infrastructure

In the rush to adopt new tools, it is easy to view ethics as a constraint – a set of brakes applied to the engine of innovation. However, the expert analysis from the e-Governance Academy on risks and ethics reframes this perspective entirely: ethics is not a constraint; it is infrastructure. Without the guardrails of fairness, accountability and transparency, the adoption of AI in public services will collapse under the weight of public mistrust.

This is particularly relevant for regions pursuing predictive analytics and automated decision support, where the expert analysis identifies a specific risk: the black box problem in funding decisions. If an entrepreneur is denied a grant based on an algorithmic score that they cannot see or understand, the legitimacy of the entire support system is undermined. Transparency, therefore, is not a feature to be added later; it is a prerequisite for the system's existence. As the expert analysis on AI ethics commissioned for this project (see Methodology on page 7) explains, for AI to be justifiable in the public sector, the rationale of a decision must be accessible to the citizen. This concept of outcome transparency requires explanations that are socially meaningful and translate technical logic into human terms.

It is therefore necessary to operationalise the FAST Track Principles (fairness, accountability, sustainability and transparency) not as abstract ideals but as engineering requirements.

FAST Track Principles

Fairness requires that the proposed recommendation engines be audited for bias. If the training data favours urban businesses, the AI will replicate that bias, further alienating the rural Traditional Artisan. The Alan Turing Institute highlights *data fairness* as a critical component, necessitating representative datasets that reflect the diversity of the population.

Accountability demands that for every automated output – whether a sector classification or a mentor match – there is a human "in the loop" or "on the loop" who bears ultimate responsibility. The state cannot outsource its duty of care to an algorithm. This aligns with the concept of *answerability* in the Alan Turing Institute framework, which requires a continuous chain of human responsibility.

Sustainability means designing systems that are resilient and maintainable. As noted in the Nortal white paper,[34] *financial sustainability* is a key strategic concern. Launching a complex AI system without the budget or skills to maintain it creates technical debt that will burden future administrations.

Transparency is the bedrock of trust. The Abundance Institute[35] points to Estonia's data tracker as a gold standard that allows citizens to see exactly who has accessed their data and why. This level of transparency transforms the user from a passive subject into an informed participant.

Without the guardrails of fairness, accountability and transparency, the adoption of AI in public services will collapse under the weight of public mistrust.

3.2 Realities of navigating the EU AI Act

The governance of AI is no longer a theoretical debate; it is a matter of hard law. The regional legal analyses (see Methodology) for Andalusia, Extremadura, Madrid, Navarra, and North Rhine-Westphalia all converge on a single, critical finding: the classification of risk under the EU AI Act.

Specifically, **Annex III, point 5(a)** of the AI Act[36] classifies AI systems used to evaluate eligibility for essential public services and benefits as high-risk. This has profound implications for the proposed use cases.

The EU AI Act establishes a risk-based classification framework. Understanding where different AI applications fall within this framework is essential for strategic planning.

34 Raieste, A., Solvak, M., Velsberg, O., McBride, K. (2025). *Government Efficiency in the Age of AI: Toward Resilient and Efficient Digital Democracies.*

35 Ilves et al. (2025). *Government Digitization Efficiency and Performance in the US and Estonia.* Abundance Institute.

36 https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3

Risk level	Requirements	Examples of entrepreneurship services
Unacceptable risk	Prohibited. Must not be deployed.	Social scoring systems; manipulative AI targeting vulnerabilities
High risk	Conformity assessment; human oversight; documentation; risk management; ongoing monitoring	AI that determines grant eligibility; automated approval or rejection of applications
Limited risk	Transparency obligations only (users must be informed they are interacting with AI)	Chatbots providing general information; AI-generated content
Minimal risk	No specific obligations (voluntary codes of conduct encouraged)	Recommendation engines suggesting resources; search optimisation; analytics dashboards

The critical distinction for the target regions is between high-risk and lower-risk categories. Most proposed AI applications can be designed to remain in the Limited or Minimal Risk categories by ensuring that they remain advisory (decision support), rather than aimed at outright decision-making. The following examples illustrate how this classification applies in practice:

- If a business portal uses AI to *determine* who receives a grant, it falls under the high-risk category, triggering onerous compliance obligations, third-party conformity assessments and strict data governance requirements.
- Similarly, if a dynamic financial support configurator automatically filters out applicants, it enters this high-risk regulatory territory.

However, the legal expert analysis (see Methodology) provides a strategic path forward through the distinction set out in this report, between *decision-making* and *decision support*. If the AI is designed to *assist* the civil servant by providing recommendations, summaries and risk scores that a human officer then reviews and approves, the regulatory burden is significantly lower. The AI remains a tool rather than a decision-maker. This aligns with the observations in the OECD report[37] regarding human–AI collaboration, in which humans use AI to filter out noise while retaining final judgment. This approach eases regulatory compliance and builds trust and *acceptance* of the technology among civil servants and citizens, thereby contributing to the development of responsible AI practice in society.

37 OECD (2025). *Governing with Artificial Intelligence.*

The governance of AI is no longer a theoretical debate; it is a matter of hard law.

This *hybrid intelligence* model is not merely a legal workaround; it is a stronger operational model. It combines the processing power of AI with the contextual judgement of the human civil servant. In practice, this means that AI can propose a business classification, but the civil servant confirms it, or that AI can suggest a personalised path for an entrepreneur, but a human adviser validates it. This approach aligns with the protections in GDPR Article 22 against solely automated decision-making, ensuring that citizens always retain the right to human intervention.

3.3 Cybersecurity and regional compliance

Trust also requires security. The integration of AI into public platforms introduces new vectors for cyber threats, from data poisoning to adversarial attacks. The regional analyses indicate that while foundational cybersecurity frameworks are in place across the target regions, they must be specifically adapted for AI. The OECD report[38] warns that AI systems can be vulnerable to "unauthorised access and data breaches", necessitating robust cybersecurity measures.

Where the protection of user or SME data is a statutory requirement, this means that any AI system must be secure by design.

This implies that data used to train models must be sanitised of personally identifiable information (PII) before it enters the model, a process known as data minimisation. The Alan Turing Institute[39] reinforces this point, emphasising the need for data fairness and integrity throughout the AI lifecycle.

Furthermore, as beneficiary regions move towards a model of public–private data sharing to support their AI ecosystems, they must establish secure, encrypted channels for this exchange. The state cannot simply open up its databases; it must curate secure data spaces in which value can be extracted without compromising citizen privacy. This is the technical foundation of the digital sovereignty concern raised in the digital transformation enablers analysis (see Methodology on page 7). If the state loses control of its data, it loses its ability to govern. The Nortal white paper[40] suggests "sovereign clouds" and a "national data commons" as architectural solutions to ensure control and resilience. This approach also aligns with the OECD recommendation[41] to build "robust digital public infrastructure" to safeguard against cyber threats.

38 OECD (2025). Governing with Artificial Intelligence.

39 Leslie, D. (2019). Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector. The Alan Turing Institute.

40 Raieste, A., Solvak, M., Velsberg, O., McBride, K. (2025). Government Efficiency in the Age of AI: Toward Resilient and Efficient Digital Democracies

41 ibid.

Chapter 4:

International reference cases

The preceding chapters identified specific technological opportunities and their associated requirements: interoperability infrastructure, data governance frameworks, digital identity systems and the human capacity to manage them.

This chapter tests the realism of those proposals against proven public-sector implementations.

This chapter tests the realism of those proposals against proven public-sector implementations. Rather than presenting these cases as models to be copied wholesale, they are examined as analytical evidence – each illuminating what is actually required to move from concept to a functioning service.

The cases were selected based on documented outcomes, transferable lessons and relevance to the constraints identified in the target regions. Each reference case is linked to one or more technology categories from Chapter 2 and to the digital building blocks identified as prerequisites in the Executive Summary. Where a case demonstrates successful implementation of a capability under consideration in the target regions, the analysis examines not only what was built but also the institutional, technical and governance conditions that made it possible.

4.1 Estonia: The networked assistant (Bürokratt)[42]

A blueprint for interoperability

Estonia's Bürokratt is often misunderstood as merely a chatbot. In reality, it is a network of interoperable AI agents. It allows citizens to ask a question of one agency and receive an answer that draws on data from several others, without the user being exposed to the complexity of the backend. This is achieved through the X-Road data exchange layer, which enables secure, authenticated communication between otherwise disparate systems.

- **Relevance.** This approach directly addresses the fragmentation crisis identified across the regions. It demonstrates that it is not necessary

42 https://www.kratid.ee/en/burokratt

to centralise all data into a single mega-database to achieve a unified user experience. Instead, an interoperable layer can be built in which distinct agents communicate. This aligns with the "federated architecture" recommended in the Digital Transformation Enablers analysis.

- **Lesson.** Start with the architecture of exchange, not just the interface. Build the plumbing before you build the tap.

4.2 Singapore: Proactive governance (GoBusiness)[43]

A blueprint for personalisation

Singapore's GoBusiness platform reverses the logic of service delivery. Instead of waiting for an SME to search for a grant, the system uses data that is already held to recommend relevant schemes proactively. It applies a "digital health check" to diagnose a business's needs before prescribing a solution. This proactive model significantly reduces the administrative burden on businesses.

- **Relevance.** This approach shifts the state from a passive repository of information to an active partner in business growth. It operationalises what a personalisation engine could look like in practice. The OECD report[44] cites Singapore's Moments of Life app as an example of proactive service delivery that anticipates user needs based on life events.
- **Lesson.** Use data to reduce the cognitive load on the entrepreneur. Shift the burden of discovery from the citizen to the state.

4.3 Ukraine: The state in a smartphone (Diia.AI)[45]

A blueprint for leapfrogging

Ukraine's experience with Diia demonstrates the power of leapfrogging. Faced with legacy corruption and inefficiency, the authorities did not attempt to fix old paper-based processes; instead, they built a mobile-first digital state from scratch. The AI assistant does not merely answer questions; it delivers services, generating official documents directly within the chat interface.

- **Relevance.** For regional administrations facing challenges related to legacy systems and rural access, this solution provides a powerful precedent. It shows that mobile-first, AI-driven interfaces can bypass the need for physical offices and desktop-based bureaucracy. The OECD[46] report notes that Diia.AI is the "world's first national AI assistant integrated into a government portal" that is capable of delivering real services.
- **Lesson.** Do not digitise the bureaucracy; digitise the output. Focus on the mobile experience as the primary interface for the citizen. In practice, this means prioritising the digital delivery of

43 https://www.gobusiness.gov.sg/
44 OECD (2025). *Governing with Artificial Intelligence.*
45 https://digitalstate.gov.ua/news/govtech/iak-ai-do-pomahaye-diyi-obrobliaty-desiatky-tysiach-zvernen-i-buty-na-zviazku-z-korystuvachamy-247-mykhaylo-fedorov-pro-vnutrishni-ai-instrumenty-komandy-pidtrymky
46 OECD (2025). Governing with Artificial Intelligence.

end services (e.g. issuing permits and licences) rather than simply placing existing procedural steps online.

4.4 Finland: Life-event service design (AuroraAI, 2020–2022)[47]

A blueprint for human-centricity

Finland's AuroraAI programme (2020–2022) organised services not by department but by "life event" (e.g. starting a business or hiring an employee). The AI operated in the background, stitching together the necessary services from various agencies to support each specific event. Although the programme has since concluded, its approach offers a valuable conceptual model for how government can rethink service delivery.

- **Relevance.** This approach addresses the silo problem identified in Germany and Spain. It requires the administration to view services from the user's perspective, breaking down artificial barriers between ministries. This mirrors the Nortal white paper's advocacy[48] for "self-organising services" that focus on capabilities rather than institutional boundaries.
- **Lesson.** Structure should follow user need, not administrative convenience. Design services around the human experience rather than the organisational chart.

47 https://oecd-opsi.org/innovations/auroraai/
48 Raieste, A., Solvak, M., Velsberg, O., McBride, K. (2025). *Government Efficiency in the Age of AI: Toward Resilient and Efficient Digital Democracies.*

Chapter 5:

Strategic recommendations

This chapter translates the findings of the innovative technologies overview into actionable, region-sensitive recommendations. It is structured to support beneficiary authorities in prioritising interventions in line with their digital maturity, institutional capacity and regulatory context.

Each recommendation is grounded in the preceding analysis and tailored to the specific needs of the five beneficiary regions: Andalusia, Extremadura, Madrid, Navarra and North Rhine-Westphalia. The recommendations are organised using a consistent framework:

- Primary constraints
- Priority shifts
- Lead intervention
- AI posture
- Regulatory strategy
- Priority AI tools
- Feasibility considerations
- Next steps

A summary matrix is provided at the end of this chapter for quick reference.

The preceding analysis reveals a fundamental insight: innovative technologies are dependent on foundational digital building blocks. Interoperability, data governance, digital identity and cybersecurity are not features to be added later but prerequisites to what is feasible – or, indeed, possible.

This principle reframes how AI adoption is approached. Governance capacity, data infrastructure and institutional readiness function as gating conditions when assessing the feasibility of any intervention. An administration cannot leapfrog these foundations. Attempting to deploy sophisticated predictive analytics on top of fragmented data systems, or to automate eligibility decisions without established accountability mechanisms, does not accelerate modernisation – it compounds existing dysfunction. Sequence matters as much as the solution.

The challenge of AI adoption in public entrepreneurship services is therefore not uniform. Each region faces a distinct configuration of constraints – institutional, technical and cultural – that determines which interventions will succeed and which will fail. Recognising this heterogeneity is not an argument against ambition; it is a precondition for effective investment.

Resources directed at interventions that are misaligned with a region's actual readiness profile risk generating frustration rather than transformation.

First, there is an institutional barrier. The fragmentation documented across the target regions – data trapped in silos, processes disjointed and registries maintained manually – is not a technical oversight awaiting a technical fix. It is the digital expression of analogue statecraft, with departmental autonomy encoded into system architecture. AI cannot resolve this fragmentation; it can only be deployed effectively once the underlying interoperability problem has been addressed. Layering intelligent tools on top of incoherent infrastructure amplifies dysfunction rather than curing it.

The preceding analysis reveals a fundamental insight: innovative technologies are dependent on foundational digital building blocks.

Second, the regions are not uniformly positioned. In some contexts, the binding constraint is digital literacy, as rural entrepreneurs cannot access sophisticated services regardless of how well designed they are. In others, the binding constraint is cultural resistance, with civil servants viewing algorithmic tools as threats to established workflows. Elsewhere, the constraint is impact data: the institutional infrastructure exists but lacks the intelligence layer required to move from reactive administration to proactive governance. A universal prescription – "deploy conversational AI" or "build a recommendation engine" – ignores these structural differences and guarantees that at least some regions will implement tools that are misaligned with their actual needs.

Third, the regulatory environment is no longer permissive. The EU AI Act's classification of eligibility-determining systems as high-risk (Annex III, point 5(a)) creates material compliance obligations that must shape system design from the outset. The strategic choice between decision support, in which AI recommends and humans decide, and decision-making, in which AI determines outcomes, is not merely philosophical; it determines whether a system triggers third-party conformity assessment, mandatory human oversight protocols and ongoing audit requirements. For most regions, remaining in low-risk advisory territory is not timidity but prudence.

The Strategic Recommendation Matrix (see Table 1) below synthesises these findings into differentiated implementation priorities. It should be read as a diagnostic tool rather than a shopping list. The first item in each column represents the binding constraint or priority shift that must be addressed before subsequent interventions can succeed. Regions that invert this sequencing – pursuing sophisticated predictive analytics before establishing basic interoperability, or automating eligibility decisions before building the user trust that legitimises algorithmic governance – will find that their investments generate friction rather than efficiency.

The recommendations in sections 5.2–5.5 elaborate principles that apply across all regions. However, their sequencing, emphasis and specific implementation must be calibrated to the regional profile

described in the matrix. What follows is not five actions that every region must take; it is five dimensions along which each region must make context-specific choices.

The matrix encodes a critical insight: the same technology serves fundamentally different functions depending on the regional context. Conversational AI in **Extremadura** is primarily an *inclusion* tool – a bridge that enables rural entrepreneurs with limited digital literacy to access state support through voice interaction. The same technology in **Germany (North Rhine-Westphalia)** serves a *trust-building* function – a non-threatening entry point for civil servants and users who are sceptical of algorithmic governance. In **Madrid**, by contrast, conversational AI functions as an *efficiency* tool, handling routine queries and freeing human advisers to focus on complex cases.

This functional differentiation has direct implications for design. An empathetic, voice-first interface optimised for the Traditional Artisan persona would frustrate Digital Explorers, who expect API access and sandbox environments. Conversely, a sophisticated data dashboard that delights advanced policymakers would be ineffective in regions where the prerequisite user base does not yet exist.

The matrix encodes
a critical insight:
the same technology serves
fundamentally different
functions depending on
the regional context.

Region	Primary constraints	Priority shifts	Lead intervention	AI posture	Regulation	Priority AI tools
Andalusia	1. Institutional fragmentation 2. Data trapped in silos 3. Poorly targeted information	1. Silos » Ecosystems 2. Compliance » Agency 3. Reactive » Proactive	Interoperability layer enabling federated search across municipal and regional portals	Decision support: AI recommends, human confirms. Focus on navigation and discovery.	Low-risk (advisory only); avoid high-risk triggers	• Federated search agents • Recommendation engine • Content curation AI • Matchmaking platform
Extremadura	1. Low digital literacy 2. System fragmentation 3. Talent shortage	1. Compliance » Agency 2. Silos » Ecosystems 3. Reactive » Proactive	Voice-activated conversational AI with empathetic guidance for rural entrepreneurs	Strictly Advisory: AI as a digital companion with human escalation built in.	Low-risk; prioritise accessibility over automation	• Conversational AI (voice) • Empathetic interface NLP • Business plan simulator • Procedural roadmap AI
Madrid	1. Fragmentation and siloed platforms 2. Bureaucratic rigidity 3. Lack of impact data	1. Reactive » Proactive 2. Silos » Ecosystems 3. Compliance » Agency	Data-driven governance dashboard with predictive analytics for policy impact	Hybrid intelligence: AI handles triage, matching and forecasting; humans validate decisions.	Accept high-risk for eligibility tools with conformity assessment	• Predictive analytics dashboard • Dynamic financial configurator • Unified search and aggregation • International talent hub AI
Navarra	1. Service fragmentation 2. Weak innovation-business links 3. Bureaucratic complexity	1. Silos » Ecosystems 2. Reactive » Proactive 3. Compliance » Agency	AI-powered matchmaking hub connecting entrepreneurs with mentors, investors and university research	Hybrid intelligence: proactive recommendations with explainable AI for trust-building.	Low-risk by default; pilot high-risk for grant matching	• Matchmaking platform • Personalisation engine • Explainable AI toolkit • Governance analytics
North Rhine-Westphalia	1. Manual registry maintenance 2. Cultural resistance to change 3. Data protection friction	1. Silos » Ecosystems 2. Compliance » Agency 3. Reactive » Proactive	AI-driven business classifier automating sector detection and business model maintenance	Decision support only: AI proposes classification, civil servant confirms. Trust is built incrementally.	Strictly low-risk; use explainable AI to address GDPR concerns	• Business classifier (NLP) • Data validation AI • Conversational assistant • Once-only integration layer

Table 1. Strategic recommendation matrix.

Reading guide: constraints and shifts are ordered by diagnostic weight (1 = most binding). AI posture reflects the decision support versus decision-making distinction under the EU AI Act. Regulatory path indicates the strategic classification approach for Annex III compliance.

5.1 Make decision support the default architecture

Across all regions, the matrix converges on a single architectural principle: AI should support decisions, not make them. This is not merely a legal accommodation to the EU AI Act's Annex III classifications; it is a stronger operational model that combines algorithmic processing power with human contextual judgement. Governments should therefore position their AI systems strategically to avoid the most onerous compliance burdens while maximising utility.

Practical implementation varies by regional capacity. In **Madrid**, where institutional sophistication permits, a hybrid intelligence model can grant AI significant autonomy over triage, matching and forecasting, with human validation at decision points. In **North Rhine-Westphalia**, where cultural resistance is the binding constraint, AI must remain strictly advisory, proposing classifications that civil servants confirm and building trust incrementally through demonstrated reliability. In **Extremadura**, human escalation must be built into every AI interaction, ensuring that the system never leaves a user without a clear path to human support.

Across all regions, the matrix converges on a single architectural principle: AI should support decisions, not make them.

The core requirement is to design systems in which AI acts as a sophisticated navigational aid rather than a final arbiter.

Decision support enablers

- **Hybrid intelligence architecture.** Governments should mandate a hybrid intelligence architecture for all high-stakes processes. This requires systems to be designed so that AI performs data analysis, pattern recognition and recommendation, while the final commitment action is reserved for a human. For example, in grant allocation, the AI should provide a risk score, a summary of eligibility criteria met and a flagged list of anomalies, while a civil servant retains responsibility for approval. This approach satisfies legal requirements for human oversight, leverages the contextual judgement of experienced staff and reduces the risk of unchecked algorithmic bias or error. It transforms the civil servant from a data processor into a decision validator.
- **Clear demarcation of autonomy.** Administrations must establish and codify clear boundaries for AI autonomy. Low-stakes tasks, such as appointment scheduling, initial information retrieval and document routing, can and should be fully automated to improve efficiency. However, high-stakes tasks involving funding decisions, legal penalties or the denial of services must remain advisory. This demarcation should be defined explicitly in internal policy and communicated transparently to users, ensuring that citizens know when they are interacting with a machine and when a human is involved.

- **Feedback loops and continuous learning.** Robust mechanisms must be implemented for human feedback on AI recommendations. A system in which a civil servant consistently overrides an AI suggestion without that information being captured is a failed system. Every override should be logged as a training signal and used to retrain the model, creating a virtuous cycle in which the system learns from human expertise and improves its accuracy and relevance over time. This transforms AI from a static tool into a learning partner.
- **Ethical design principles.** AI systems must align with core democratic values, such as fairness, accountability and transparency. This includes incorporating bias mitigation strategies during data collection and model training, as well as the provision of clear explanations for AI-generated decisions to maintain public trust.

5.2 Establish the integration layer as a priority

For **Andalusia**, **Navarra** and **North Rhine-Westphalia**, where "Silos » Ecosystems" ranks as the first or second priority shift, no AI application will deliver value until the underlying fragmentation is addressed. The strategic priority is therefore the establishment of an interoperability layer before deploying user-facing AI tools.

This does not require wholesale system replacement. Following established interoperability models, such as Estonia's X-Road, the Netherlands' NLX and similar national data exchange frameworks, the goal is a "data broker" layer that enables secure, authenticated communication between existing systems. In practice, this means federated search across portals, automating data flows that currently require manual registry maintenance, and connecting fragmented landscapes of development agencies into a coherent user experience.

Some regions present an exception that proves the rule. Where the binding constraint is not system fragmentation but user access, with digital literacy so low that sophisticated back-end integration would be invisible to the primary user base, the priority is reversed: build the accessible front end first, such as voice-activated AI and empathetic interfaces, and then invest in integration as user adoption grows.

The allure of deploying advanced AI tools often distracts from the unglamorous but essential work of infrastructure. The regional analyses consistently identify data fragmentation as a primary barrier. Therefore, the strategic priority must be the establishment of an integration layer. Before procuring AI solutions, the data pathways must be put in place.

Secure data exchange enablers

- **Data brokerage as infrastructure.** Develop a centralised data broker layer or interoperability platform, drawing on models such as Estonia's X-Road or similar national data exchange frameworks, that acts as a secure intermediary between disparate agency databases. This layer should handle data requests, authentication and translation between different data formats. It enables the seamless flow

of information without requiring a large-scale, high-risk overhaul of legacy systems. This is the technical realisation of the Once-Only Principle, ensuring that citizens do not have to act as couriers for their own data between government departments.

Before procuring AI solutions, the data pathways must be put in place.

- **API standardisation mandate.** Enforce a strict API-first policy across all government agencies. Any new digital service or database that is procured or developed must expose open, documented APIs that enable secure data exchange. This prevents the creation of new data silos and ensures that the ecosystem remains interoperable by default. It shifts procurement away from monolithic software suites towards modular, connectable components. This approach aligns with broader EU interoperability guidance, including the European Interoperability Framework, and reflects recognised best practice in digital government.
- **Targeted data quality campaigns.** Launch targeted data quality improvement programmes focused on the high-value datasets required for AI. AI models are only as reliable as the data on which they are trained. Investment is required in cleaning, standardising and tagging key datasets, such as business registries and tax records. This is a critical prerequisite for reliable predictive analytics. Poor-quality data will undermine even the most sophisticated AI models.
- **Secure data sharing protocols.** Establish robust protocols for data sharing between the public and private sectors, ensuring compliance with data protection regulations, such as GDPR. This includes the use of encryption, anonymisation and access controls to safeguard sensitive information while enabling responsible data exchange to support AI innovation.

5.3 Adopt a persona-centric design philosophy

Service design must move beyond generic user interfaces to address specific entrepreneurial needs. Analysis of user needs shows that a one-size-fits-all approach will inevitably fail to serve the diversity of the entrepreneurial ecosystem. The strategy must therefore focus on adaptive interfaces that meet users where they are.

Persona-centric design enablers

- **Multimodal interfaces for inclusion.** Develop service interfaces that support multiple modes of interaction. While traditional web forms remain necessary for complex submissions, they are a barrier for many users. Mobile applications should be prioritised for on-the-go access. Crucially, for rural and less digitally literate users, voice-activated assistants are essential. These interfaces should allow users to navigate complex bureaucracy using natural language, effectively leapfrogging the need for deep digital literacy.

- **Dynamic, tailored user journeys.** Use AI to generate personalised user journeys dynamically based on the entrepreneur's profile. The system should detect who the user is – whether a rural artisan or a technology startup founder – and adapt the interface accordingly. The artisan should encounter a simplified, visual interface focused on grants and local support, while the technology founder should be presented with API documentation, sandbox environments and investment opportunities. The burden of navigation should be shifted from the user to the system.
- **Proactive outreach.** Shift the fundamental model of service delivery from a reactive pull to a proactive push. Use data analytics to identify entrepreneurs who may be eligible for specific programmes or support measures and reach out to them proactively with personalised offers. This reduces the cognitive load on users – for example, uncertainty about available grants – and ensures that valuable support reaches those who need it most, rather than only those with the resources to find it.
- **User feedback integration.** Continuously collect and analyse user feedback to improve service design iteratively. This includes using sentiment analysis and usage data to identify pain points and optimise the user experience, ensuring that digital services remain relevant and user-friendly.

5.4 Implement accountability by design

Trust is the currency of the digital state. To counter the black box problem and public scepticism, transparency mechanisms must be embedded in the technical architecture of AI systems from day one. Accountability by design means that systems are built to be auditable.

Accountability-by-design enablers

- **Mandatory explainable AI (XAI).** Require the use of XAI techniques for all public-facing AI systems. Users should be able to access a plain-language explanation of why a specific recommendation was made or why a particular score was assigned. This demystifies the technology and builds confidence in its fairness. "Because the algorithm said so" is not an acceptable justification in a democratic context.
- **Immutable audit trails.** Ensure that every automated recommendation or decision is logged in a secure, immutable audit trail. This trail should record data inputs, the model version used, the output generated and any human intervention. This is essential for post hoc analysis, debugging and accountability in the event of a dispute. It protects both the citizen and the civil servant.
- **Public algorithm registers.** Publish a comprehensive register of all algorithms used in public service delivery, detailing their purpose, data sources, logic and ownership. This level of transparency enables public scrutiny and helps dispel concerns about "secret

algorithms" shaping decisions that affect citizens' lives. It signals that the state has nothing to hide.

- **Regular algorithmic audits.** Conduct regular audits of AI algorithms to identify and mitigate bias, errors or unintended consequences. These audits should be performed by independent third parties to ensure objectivity and to build public trust in the fairness and reliability of automated systems.

5.5 Define a regulatory strategy by region

The EU AI Act's classification of eligibility-determining systems as high-risk (Annex III, point 5(a)) creates a strategic choice for each region. The matrix encodes differentiated approaches based on institutional capacity and risk tolerance.

Madrid, with its robust governance infrastructure and policy ambitions, can accept high-risk classification for tools such as the Dynamic Financial Support Configurator, bearing the conformity assessment burden in exchange for fuller automation of eligibility determination. **Navarra** may pilot high-risk classification for grant matching while keeping other systems in low-risk advisory mode.

For **North Rhine-Westphalia**, **Andalusia** and **Extremadura**, the strategic imperative is to remain in low-risk territory. This requires designing AI systems that recommend rather than determine outcomes, by providing risk scores, eligibility summaries and flagged anomalies that human officers review and approve. The regulatory burden is lower and, critically, this approach aligns with the trust-building and accessibility priorities these regions require.

5.6 Build technical sovereignty as a long-term horizon

All regions share a common long-term imperative: building internal technical capacity to govern, manage and evolve their AI infrastructure. Immediate priorities differ – some regions require basic digital skills programmes, while others need specialist AI governance expertise – but the direction of travel is uniform.

This requires the creation of senior data governance and strategy roles, such as Chief Data Officers and Chief AI Officers, with genuine authority over digital transformation and data policy. It requires continuous upskilling programmes calibrated to regional starting points. It also requires the adoption of "public money, public code" principles where feasible, prioritising open-source solutions that prevent vendor lock-in and enable cross-regional collaboration.

Technical sovereignty means possessing the internal capability to govern, manage and evolve digital infrastructure. The state cannot be an effective purchaser of AI if it does not understand what it is buying, when to collaborate and how to hold vendors accountable when systems fail to deliver.

Technical sovereignty enablers

- **Technical leadership roles.** Create and empower senior technical roles within the civil service. These roles must carry the authority and resources required to drive digital transformation and to challenge vendors where necessary. They should bridge the gap

between policy and technology, ensuring that technical decisions align with strategic public goals.

- **Continuous learning and upskilling.** Implement robust, continuous learning programmes for civil servants at all levels. This goes beyond basic digital literacy. Frontline staff require training in working *with* AI tools. IT professionals need specialist training in managing AI systems. Leadership requires strategic training in AI governance and ethics. A knowledgeable workforce is the most effective safeguard against vendor capture and project failure.
- **Public money, public code.** Adopt a "public money, public code" policy, prioritising open-source solutions where feasible. This approach prevents vendor lock-in, allows for greater customisation and enables collaboration with other governments facing similar challenges. It ensures that the state retains ownership of the code that runs its critical infrastructure, allowing for continuous improvement and adaptation without dependence on proprietary licences.
- **Partnerships.** Foster strategic partnerships with academia, research institutions and the private sector to leverage external expertise while retaining control over core digital assets. Such collaborations can accelerate innovation and knowledge transfer, strengthening the public sector's technical capabilities.

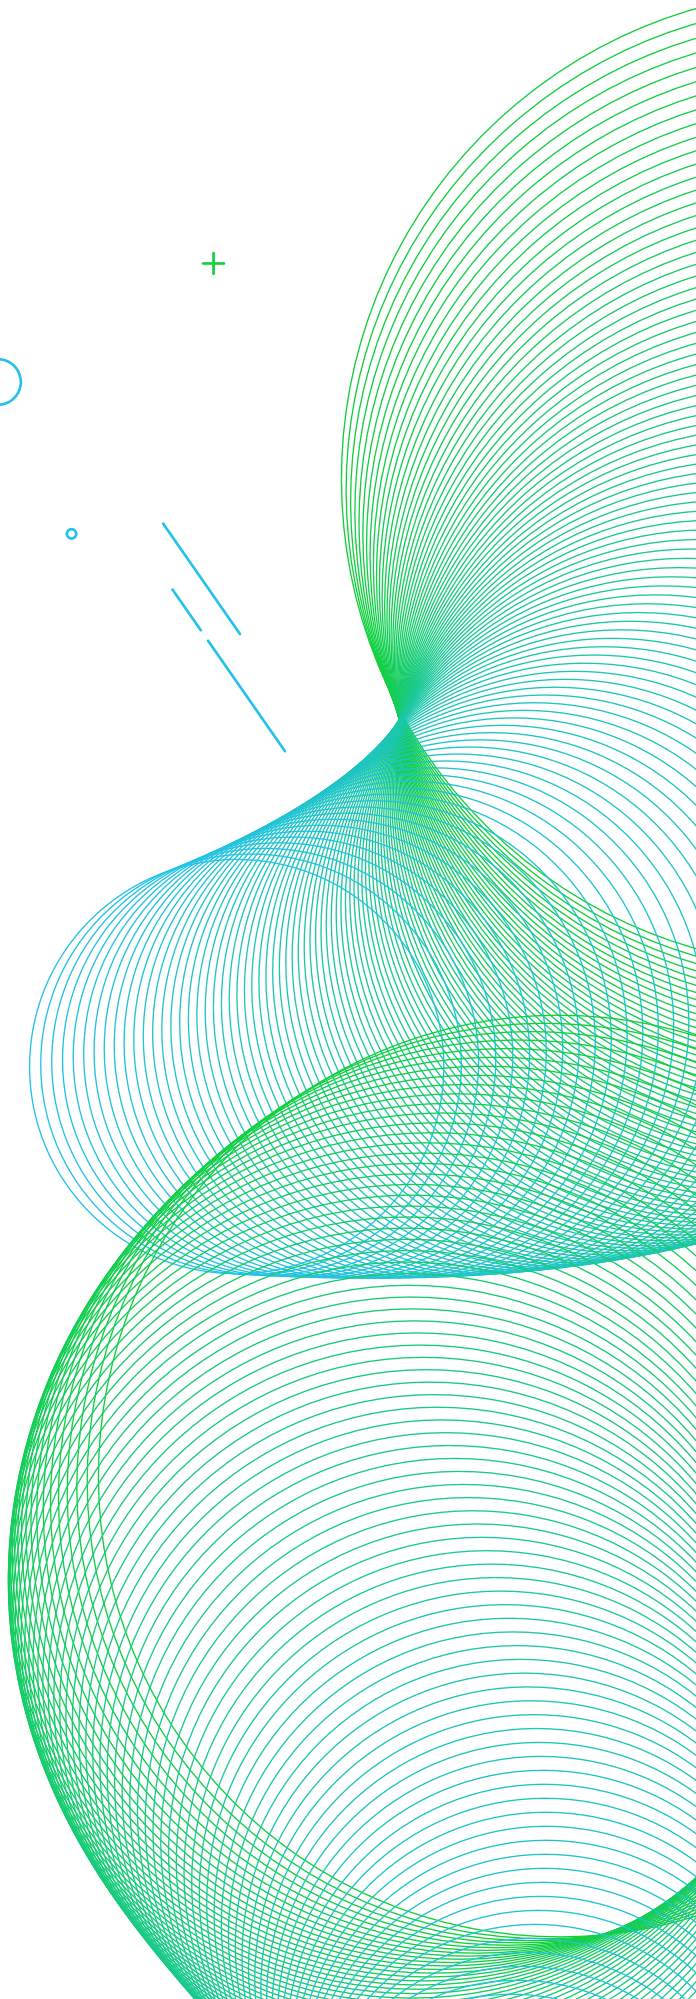

Conclusion

This report began with a diagnosis: the primary barrier to innovation in public entrepreneurship services is not a scarcity of technology but a legacy of institutional design. Across the regions analysed, entrepreneurs encounter fragmented registries, disconnected portals and administrative silos that force them to act as couriers of their own data between government departments. The state, unable to see the entrepreneur as a whole entity, delivers friction rather than support. Digitisation efforts that merely replicate paper processes on screens have failed to address this structural dysfunction. The challenge, as argued, is not to digitise the existing state but to redesign it.

The challenge is not to digitise the existing state but to redesign it.

Layering AI onto these fragmented foundations will not resolve the dysfunction. The report's central argument is that sequencing matters: interoperability before intelligence, data quality before prediction and accountability before automation. Regions that respect these dependencies have a clear path forward. Those that skip steps – rushing to deploy chatbots on top of siloed registries, or predictive models trained on incomplete data – risk investing resources in systems that cannot deliver on their promise.

At the same time, the scale of the opportunity is equally clear. The state has access to information that, if mobilised coherently, could transform the entrepreneurial experience by identifying relevant permits before the entrepreneur asks, surfacing relevant funding before deadlines pass and connecting businesses to support networks based on their actual profile rather than their ability to navigate bureaucracy.

This is not speculative. The foundational technologies exist. The legal frameworks, including the EU AI Act, provide guardrails that, when properly understood, enable rather than obstruct innovation. What is required is institutional commitment to build the prerequisites: interoperability layers, data governance frameworks and the internal capacity to manage and evolve these systems over time. The vision is a state that anticipates needs rather than reacting to requests, recognises the entrepreneur as a whole person rather than a sequence of transactions and earns trust through transparency.

The following recommendations translate these insights into action. They do not offer a universal prescription but rather a framework for choices that must be calibrated to each region's starting point and constraints.

Cross-cutting recommendations

Build interoperability first

This is the gating condition. No AI system will succeed without a foundational data exchange infrastructure. Prioritise API-first policies and federated data layers before investing in advanced applications.

Prioritise data quality

AI is only as good as the data it learns from. Launch targeted data-quality programmes focused on high-value datasets, including cleaning, standardising and tagging business registries, tax records and service usage data. A disorderly dataset renders even the most sophisticated AI model useless.

Opt for decision support (skip decision-making for now)

To comply with the EU AI Act and build public trust, AI should assist human officials rather than replace their judgement. Systems that remain advisory avoid high-risk classification and maintain clear lines of accountability.

Design for personas

Use identifiable, context-dependent archetypes to guide service design and interface choices. A one-size-fits-all approach will inevitably fail at least one group.

Embed ethics and transparency

Operationalise the FAST Track Principles (fairness, accountability, sustainability, transparency) in all AI systems. This is not a constraint on innovation; it is the infrastructure that makes public trust possible.

Invest in capacity

Launch targeted upskilling programmes for civil servants to manage, validate and govern AI tools. Without internal expertise, administrations become consumers of technology rather than architects of public value.

Plan for sustainability

Ensure the financial and technical sustainability of AI systems from the outset. Avoid vendor lock-in by prioritising open standards, modular procurement and "public money, public code" principles where feasible.

Calibrate to context

Use the Strategic Recommendation Matrix to match interventions to regional readiness. The same technology serves fundamentally different functions depending on local constraints; what works as an accessibility bridge in one region may function as an efficiency tool in another.

Authors

Federico Plantera	Researcher, Centre for European Policy Studies (CEPS) and Associate Researcher, e-Governance Academy. Works on tech policy and AI governance projects, with research on institutional technology adoption and the political economy of AI in public administration.
Piret Hirv	Senior Expert and Head of Data Management, e-Governance Academy. Focuses on the digital transformation of public services, with extensive experience in e-governance strategy, capacity building and cross-border digital service delivery across European public administrations.

Contributors

Andres Vimb	Senior Expert, e-Governance Academy. Experience across the public and private sectors, with expertise in tax system modernisation and information systems integration.
Kalle Tomingas	PhD, Chief Data Scientist, Orion Governance. Specialist in enterprise information management, data governance and technology scalability assessment.
Konstantyn Krupskyi	Senior Expert, e-Governance Academy. Lead business analyst with over 10 years' experience in public e-services, including on Ukraine's Diia mobile application since its launch.
Diletta Huyskes	PhD, CEO and Co-founder of Immanence. Researcher in AI ethics specialising in automated decision-making in public administration and algorithmic impact assessment.
Mykhailo Kornieiev	Senior Expert, e-Governance Academy. Expert in digital transformation and open data with hands-on experience in Ukraine's e-governance sector since 2018.
Viljar Peep	PhD, Senior Expert. Former Director General of the Estonian Data Protection Inspectorate (2008–2018), with expertise in privacy law, electronic signatures and digital business registration.
Radu Serrano	PhD, Expert, e-Governance Academy Cybersecurity Competence Centre. Project Manager for the National Cyber Security Index (NCSI), with research on cybersecurity governance and policy.

e-Governance Academy
Ahtri 6, 10151 Tallinn, Estonia
+372 663 1500 | info@ega.ee | ega.ee
Facebook, LinkedIn: egovacademy