Deliverable 3.3: Policy Recommendation Toolkit V1

29-04-2024

Version 1.0

Funded by the European Union	*Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Agency. Neither the European Union nor the granting authority can be held responsible for them.*

PROPERTIES	
Dissemination level	Public
Version	1.0
Status	Final
Beneficiary	UBI
License	This work is licensed under a Creative Commons Attribution-No Derivatives 4.0 International License (CC BY-ND 4.0). See: https://creativecommons.org/licenses/by-nd/4.0/

AUTHORS		
	Name	**Organisation**
Document leader	Ntalaperas Dimitris	UBI
Participants	Xanthi Papageorgiou	UBI
	Nikos Kalatzis	UBI
	Konstantinos Tzelaptsis	UBI
	Septimiu Nechifor	SIE
	Raluca Maria Repanovici	SIE
	Iuliana Stroia-Vlad	SIE
Reviewers	Dimitris Kotios	UPRC
	Silvina Pezzetta	WLC

VERSION HISTORY				
Version	**Date**	**Author**	**Organisation**	**Description**
0.1	01/03/2024	Ntalaperas Dimitris	UBI	ToC
0.3	10/10/2024	Ntalaperas Dimitris, Xanthi Papageorgiou	UBI	Requirements
0.5	01/04/2042	Ntalaperas Dimitris, Xanthi Papageorgiou, Nikos Kalatzis, Septimiu Nechifor, Raluca Maria Repanovici, Iuliana Stroia-Vlad	UBI, SIE	Architecture
0.7	12/04/2024	Ntalaperas Dimitris, Nikos Kalatzis, Xanthi Papageorgiou, Konstantinos Tzelaptsis, Septimiu Nechifor, Raluca Maria Repanovici, Iuliana Stroia-Vlad	UBI, SIE	Prototype description
0.75	13/04/2024	Ntalaperas Dimitris	UBI	Editing
0.8	22/04/2024	Ntalaperas Dimitris, Xanthi Papageorgiou	UBI	1st draft for internal review
1.0	29/04/2024	Ntalaperas Dimitris, Xanthi Papageorgiou	Ubitech	Final draft addressing internal review comments

Table of Contents

List of figures

List of Tables

Abbreviations

Abbreviation	Description
DAO	Decentralized Autonomous Organization
dApp	decentralized (decentralised) Applications
DSU	Data Sharing Unit
EBSI	European Blockchain Services Infrastructure
eIDAS	electronic Identification and Trust Services
ESSIF	European Self Sovereign Identity Framework
EVM	Ethereum Virtual Machine
GDPR	General Data Protection Regulation
HLF	HyperLedger Fabric
SSI	Self-Sovereign Identity

Abstract

The present document presents the first iteration of the Policy Recommendation Toolkit (PRT), which aims to facilitate organizations in policy-making and provide a transparent governance model that gives the opportunity to citizens to audit these processes of policy-making and to actively participate in the formation of them via blockchain-enabled co-creation. In this first iteration, the requirements derived from the pilots are gathered; though the requirement engineering process follows an agile methodology and the requirements can thus be refined in future iterations, the core aspects of the system are identified at this stage. These aspects are consolidated in the PRT architecture, which is another outcome of the present deliverable. Implementation of the PRT and integration into both the Visualization Workbench and the decentralized infrastructure has already started; the status of the prototypes is given in the present report.

1 Introduction

1.1 Purpose and scope

The present document offers the 1st iteration of the Policy Recommendation Toolkit (PRT). Following the user stories, it derives the first set of user requirements and, based on these requirements, presents the architecture that fulfils them. At this stage and in accordance also with the methodology followed in D3.1, the first version of the architecture is limited in the business and application layers; the technical layer (application and infrastructure) will be presented in the 2nd and final iteration of the document when the finalization of the requirement engineering process and the iterative implementation process can produce the appropriate definition of the appropriate technical Solution Building Blocks (SBBs).

Although the pilot use cases focus on data and actors from organizations, the present document extends, whenever applicable, the use cases to citizens and have them to participate in the policy-making process. Following the model of open democracy and recognizing that policies primarily affect the citizens, we leverage the decentralized infrastructure defined in D3.1 to facilitate the participation of citizens in the policy-making process. This can happen via various means that are implemented via self-governed smart contracts, which also allow citizens to form opinions, provide feedback to policymakers and vote for recommended policies. These mechanisms are under investigation and implementation, with the present document providing an overview of the technology enablers that complement those proposed in the Decentralized Data Governance Framework (D3.1) and are specific to the citizen user group.

Lastly, the current status of implementation is described, along with the next steps that are planned for the following period.

1.2 Document structure

The present document is structured as follows:

- Section 1 contains the present introduction.
- Section 2 lists the requirements and briefly presents the new technology enablers that are going to be leveraged for implementing the PRT.
- Section 3 presents the various layers of the architecture.
- Section 4 presents the status of development and lists the next steps toward the finalization of the toolkit.
- Section 5 gives the conclusions of the present work.

2 Use Cases and requirements.

A policy-making process can be defined as a collaborative process that involves interest groups and analytical frameworks with the goal of forming a common set of goals and actions (Thatcher et al., 2015). To facilitate the optimal design of policies, it is essential that large groups of affected stakeholders are able to form networks in order to communicate ideas and needs and form policies using a co-creation process so that all stakeholder interests are imprinted in the resulting policies. In the domain of public policies, in particular, this process needs not only to involve large segments of the public sector but also to be transparent to the public; the citizens should be able both to co-design policies and audit them.

Modern trends in digitalization and AI can help these processes grow and enforce transparency in various facets of their execution. Although the facets of the processes that can be enhanced by digitalization are interconnected, we can roughly separate them into the following categories.

- Semantic interoperability: Policies often depend on terminology, models, and datasets that are used and understood at different levels by interested stakeholders. Semantic interoperability mechanisms ensure that when policies are defined, their constituents have a specific meaning that is understood unambiguously by all interested parties.
- Promotion of inclusiveness, responsiveness and accountability by enabling the model of Open Democracy (Landemore, 2020). Technological enablers such as the blockchain technology can help institutions and citizens to participate in activities of policy formation in its various phases, from consultation to voting. Decentralization can also act vertically through all aspects of digitalization by enforcing trust.
- Recommender systems can help produce optimal policies by solving optimization problems and suggesting candidate policies according to the constraints set, greatly reducing the complexity of designing a policy from scratch. With the advent of AI, these optimization processes can produce solutions that are very close to global optima; moreover, by using generative AI, new areas and potential actions can be explored by searching through the available datasets.

In the following sub-sections the following axes will be investigated, together with a mapping that shows how they can be leveraged to enhance policy-making in the AI4Gov Pilot Cases.

2.1 Semantic Interoperability

From the first days of the Semantic Web (Semantic Web – W3C), the goal of semantic interoperability is to provide unambiguous meaning to data exchanged between information systems. In practice, this can be very difficult since these meanings depend on context and are often shared between systems and processes that were not designed initially to be working together. The word “safe” for example, can have a different meaning depending on the domain (e.g., mean an acceptable level of emission in the domain of green growth or applied in the working conditions in labor policies). When, as it is commonly the case, policies involve multiple domains and different datasets are combined, the categorical labels of data that have a semantic relation, have to be grouped together. While this process can be done manually by a data curator, it is often tedious and, in the case of large datasets, could prove impossible. Novel techniques that involve Knowledge Graphs and AI, however, can be used to cluster entities that are semantically interlinked.

For the use cases of AI4Gov, the requirements for semantic interoperability can be seen in Table 1. We distinguish two cases:

- The pilots in isolation. This is represented in the first three rows of the table. For this case, the main requirement is to uplift the data to a taxonomy so that it can be readily interpreted and consumed by a 3rd party.
- Collaboration between pilots, in the two last rows. This is the case in which water management policy design (DPB) and waste management policy design (VVV) can be benefit by policy data maintained by JSI. How this is done is investigated in the AI Policy Making Section, however it requires that data from the two sources are aligned via an appropriate model.

Table 1: Semantic Interoperability in AI4Gov Use Cases

<table>
<tr><th>Use Case</th><th>Pilot</th><th>Data</th><th>Users</th><th>Requirements</th></tr>
<tr><td>Water Management – drinking water</td><td rowspan="2">DPB</td><td rowspan="2">-Sewage Treatment data
-Water cycling billing data
-Streaming sensor data</td><td rowspan="2">-Workers at the municipal consortium for water management
-Local administration</td><td rowspan="2">-Uplift data</td></tr>
<tr><td>Water Management – sewage water</td></tr>
<tr><td>IRCAI global 100 projects</td><td rowspan="3">JSI</td><td rowspan="3">-IRCAI data of projects submitted (textual description, URLs)
-Event Registry data (news and event items)</td><td rowspan="3">-Teams in private or public Institutions/Organizations that are submitting projects to the IRCAI Global Top 100 program.
-Government
-Corporate
-Researchers</td><td rowspan="3">-Uplift report metadata</td></tr>
<tr><td>SDG Observatory</td></tr>
<tr><td>OECD policy document analysis</td></tr>
</table>

		-OECD AI policy initiatives		
Parking tickets monitoring	VVV	-Census data -Household water data -Tourist data (arrivals, overnight stay, cruise data) -Airport traffic data -Municipality events attendance data	Policy makers	Uplift data
Waste management – Pay as you Throw				
Water Management policies	DBP-JSI	Data from DBP and JSI	Policy makers	Align data
Waste Management policies	VVV-JSI	Data from DBP and JSI	Policy makers	Align data

2.2 Open Democracy

While, in theory, citizens have access to public information, can monitor government and participate in public consultation, participation in these processes is often hindered in practice. Citizens often have to actively search for the appropriate channels, while efficient participation in the public consultation may be poisoned by deep fake, paid digital accounts and bots.

Blockchain is a technology enabler that can facilitate inclusive democratic processes while helping avoid the aforementioned caveats. The blockchain enabler and underlying infrastructure that is going to be used in AI4Gov to implement a fully decentralized data governance framework has been documented in D3.1. In this section, we are describing how the decentralized infrastructure, along with the usage of smart contracts and the on-chain data governance framework, can be used to activate citizens in policy making. The requirements for this are listed in Table 2. Briefly, the affected actors can be separated into two categories:

- Policymakers (as members of public institutions) propose policies and define the governance policies by which the policies can be endorsed.
- Citizens can vote on the proposed policies and can demand to retrieve explainability reports if the policies are based on the output of an AI algorithm.

In case where an AI output is produced deterministically, it can also run on the blockchain as a smart contract. In this case, the citizens have an extra tool for auditing, as they can validate the models by rerunning them on their nodes.

The cases have been limited to two pilots, namely DPB and VVV, since these pilots directly involve policymaking that has the potential to engage stakeholders from the whole spectrum, from public servants to citizens. The special user group "Governing body" is reserved for the users that have the right to alter the policies by which the blockchain is governed (e.g., change the voting system from unanimous to majority, change the definition of the code running a recommender system, etc.). The physical users to which this group corresponds may vary depending on the use case. For the DPB and VVV, these would be the policy makers (the users of the institutions). However, this can change if more open policy governance models are to be followed in the future. While this group is not listed in the architecture, we list it here separately in Table 2 to highlight this important distinction.

Table 2: Digital Open Democracy in AI4Gov Use Cases

Use Case	Pilot	Data	Users	Requirements
Water Management – drinking water	DPB	-Sewage Treatment data -Water cycling billing data	-Policy Maker	-Propose policy -Alter policy -Vote policy
			-Governing body	-Define governance model
Water Management – sewage water			-Citizens	-Vote policy -Audit AI

Parking tickets monitoring	VVV	-Census data -Household water data -Tourist data (arrivals, overnight stay, cruise data) -Airport traffic data -Municipality events attendance data	-Policy makers	-Propose policy -Alter policy -Vote policy
			Governing body	-Define governance model
Waste management – Pay as you Throw			-Citizens	-Propose policy -Alter policy -Vote policy

2.2.1 Trust in an open democracy

One of the key challenges that any digital platform that tries to implement mechanisms of open democracy has to face is that of trust. Citizens will avoid entering an open and inclusive platform if it lacks the appropriate transparency and trust mechanisms. Any such platform should guarantee that:

- Any feedback and consultation that is signed by the citizen cannot be altered or isolated in any way.
- Both the governance process and the mechanisms by which this process can change are clear to citizens.
- Any piece of evidence (e.g., OECD report) that is used for forming a policy can be retrieved and inspected by citizens.
- Secrecy of vote should be possible.

The first three elements are typical use cases of a blockchain infrastructure: the technological enablers and the mechanisms for enforcing them are documented in D3.1. The secrecy of the vote, however, is a new requirement that is specific to the PRT and has to be treated separately.

At first sight, providing vote secrecy seems incompatible with the nature of the blockchain; ballots need to be counted by the smart contract that implements the voting mechanism and have to be recorded into the blockchain in plain view of all peers. However, by following certain cryptographic protocols, such as the Zero Knowledge Proof (ZKP) (Feige et al., 1987), this can be used to prove certain statements without disclosing further information.

A graphical way of understanding the basic idea behind ZKP is the story of the Ali Baba cave. The setting is depicted in Figure 1. The cave has a door that connects paths A and B. The door has a code that the Prover knows. She wants to prove to the Verifier that, indeed, she knows the code without disclosing the code itself.

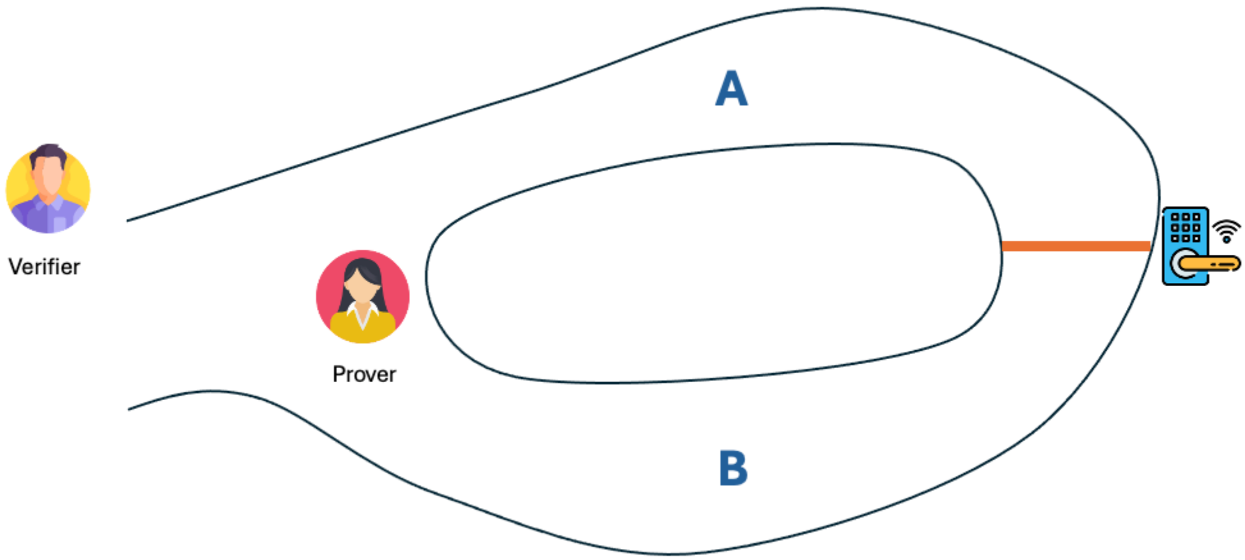

Figure 1: The Ali-Cave. The Prover knows the combination of the lock that is deep in the cave. She wants to prove to the Verifier that she knows the code without disclosing it.

A straightforward way to achieve this is to have the Prover and Verifier both randomly choose a path. First, the Prover follows the path to reach the door without the Verifier seeing, and then the Verifier goes to the entrance and shouts his choice. The Prover then has to follow the path that the Verifier called and appear in the corresponding entrance. In case the prover does not know the password, she cannot cross the door and must, therefore, return by the way she took. This path has a 50% probability of coinciding with the path that the Verifier called. If she knows the password however, she can appear on the called path. By repeating the experiment enough number of times to eliminate the chance of luck, the Prover can prove to the Verifier that she, indeed, knows the password to the door.

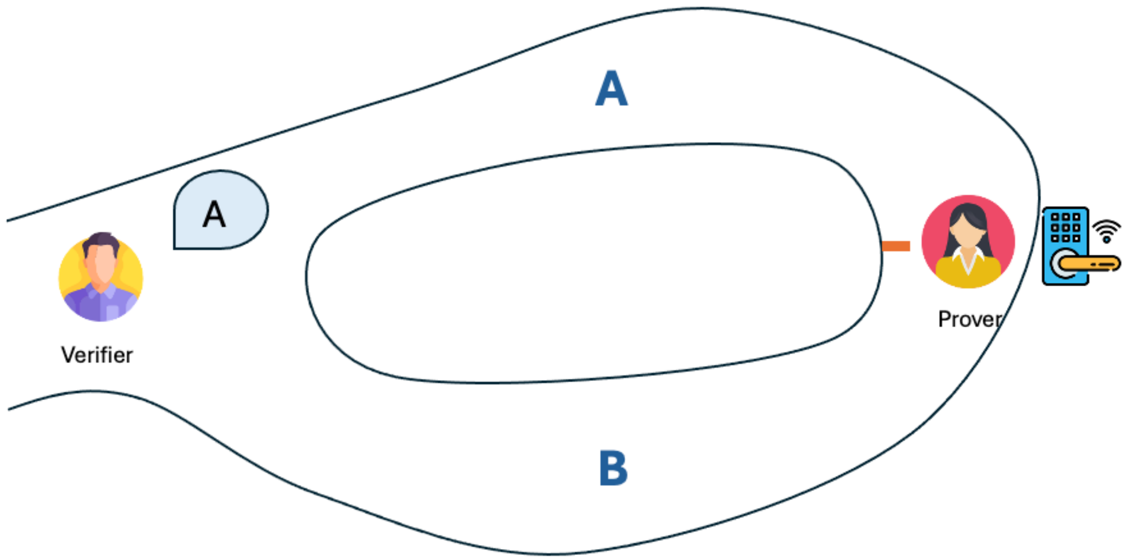

Figure 2: The Prover follows a path and opens the door. The Verifier shouts a random path. The Prover is expected to appear on the entrance corresponding to the path that the Verifier called.

Extending the above reasoning to the voting system, it can be seen that the information that is needed for deciding the outcome of a vote is not the individual ballots but rather:

- The aggregates of all the ballots
- The knowledge that a voter has cast a ballot to avoid double voting.

ZKP mechanisms can be applied to prove that a voter has cast a ballot. For computing the aggregates, various schemes based on ZKP exist, such as the "Commitment Scheme". For the purposes of the PRT however, the most promising solution that is now under development is that based on the Homomorphic Encryption. Homomorphic Encryption is a technique that allows operations on encrypted data without the need to decrypt it. The main idea of the mechanism is depicted in Figure 3. A plaintext *m* can be encrypted in the cipher *c(m)*, and consequently, if *f* is any function, the message f(m) is encrypted into the cipher *c(f(m))*. If the encryption is such that by applying *f* to *c(m)* we get the same cipher *c(f(m))* as that we would get if we encrypted *f(m)* directly, then the encryption scheme is a homomorphic encryption that allows computation of function *f* directly on the encrypted data.

As an example of this, consider the Pallier function defined by:

$$C(m) = g^m r^n mod n^2$$

with *g,n* being the public key and *r* a random number.

then

$$\begin{aligned} C(m_1)C(m_2) &= g^{m_1} {r_1}^n mod n^2 \; g^{m_2} {r_2}^n mod n^2 \\ &= g^{(m_1+m_2)} (r_1 r_2)^n mod n^2 \\ &= C(m_1 + m_2) \end{aligned}$$

It can be seen the cipher contains the sum of the encrypted sum of the messages; this is exactly what is required by a voting system.

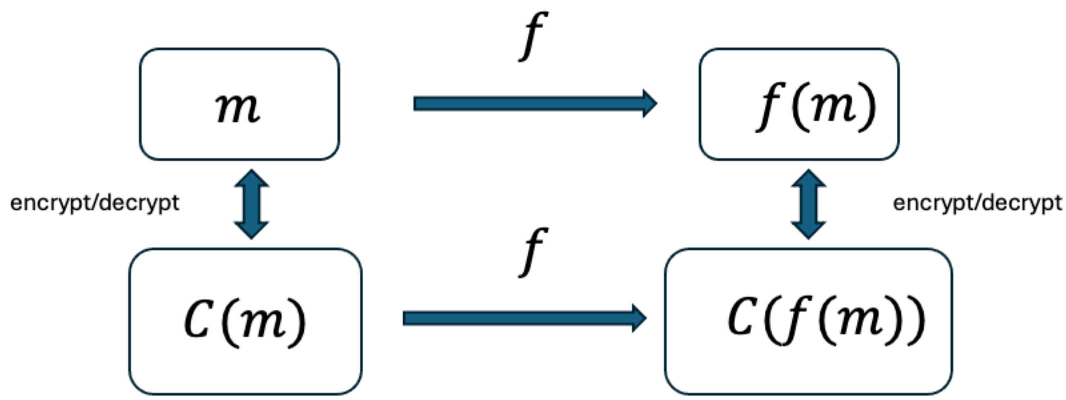

Figure 3: Homomorphic encryption. Applying f directly to c(m) produces the output c(f(m)). Decrypting it, we get the same output as we would get if we applied the function directly in the plaintext data.

2.3 AI In Policy Making

As the number of datasets that correspond a) to the underlying domain(s) of the policy under consideration and b) to the number of opinions formed in public discourse keeps growing, processing of this information via AI, both traditional and generative can lead to recommendations of new policies that benefit from insight gained by these data, that is hard to get via traditional means. AI, in this sense, can lead to breakthroughs in policy-making. However, some caveats can identified:

- The datasets may be poisoned by errors and/or bias.
- In the current state of social media, many of the data points have been themselves been generated using generative AI; these data points may too poison the AI models.
- Citizens and representatives cannot verify the source or the validity of the data; even worse, they cannot identify if the AI has been trained on such poisoned datasets and if, thus, can be trusted.

As with open democracy, decentralization can help in enforcing trust in AI models by demanding that:

- Each AI is assigned a decentralized identifier (DID), which is attached to any report it generates
- Any AI derived result or report is anchored in the blockchain together with the metadata of the AI that produced it, so that stakeholders can audit the AI to retrieve information such as its algorithm, its training parameters, its learning corpus etc.

In addition, these caveats are addressed by complying whit GDPR and the recently passed AI Act. As an extra layer of trust, AI that produces deterministic results can run on-chain as a smart contract. The execution of an AI that is implemented as smart contract is fully reproducible by any peer; in this sense stakeholders can validate the results independently. The requirements that codify the above considerations can be viewed in Table 3.

Table 3: AI-enabled Policy Making in AI4Gov

Use Case	Pilot	Data	Users	Requirements
Water Management – drinking water	DPB	-Sewage Treatment data -Water cycling billing data	Policy Makers	-Set criteria -Get recommended policies -Generate new policies
Water Management – sewage water			Citizens	-Audit AI -Get Explainability report
Parking tickets monitoring	VVV	-Census data -Household water data -Tourist data (arrivals, overnight stay, cruise data) -Airport traffic data -Municipality	-Policy makers	-Set criteria -Get recommended policies -Generate new policies
Waste management – Pay as you Throw			Citizens	-Audit AI -Get Explainability report

		events attendance data		
Water Management policies	DBP-JSI	Data from DBP and JSI	-Policy makers	Get relevant reports
Waste Management policies	VVV-JSI	Data from DBP and JSI	-Policy makers -Citizens	Get relevant reports

3 Architecture

Following the requirements laid out in Section 2, the architecture of the Policy Recommendation Toolkit can be derived. This architecture describes the toolkit to a level of granularity that is at a lower level than the one described in the AI4Gov Reference architecture that was described in D23 (Figure 4). The red boxes indicate the backend of the PRT, while the green box corresponds to its backend. Elements of the PRT overlap with the Decentralized Data Framework, which was described in D3.1, while the front end is tightly integrated with the Visualization Workbench. As such, certain elements of the architecture of the PRT will refer to elements described in D3.1; these overlaps will be dully noted.

For a uniform presentation, the same approach that was followed in D3.1 to describe the architecture will be followed here. The Archimate modeling language (Archi – Open Source ArchiMate Modelling) of The Open Group will be followed and the architecture will be described in the Business and Application Layer; as there is a strong semantic component to the PRT, a special Semantic View will also be given. As is the case with the Decentralized Data Framework, further decompositions of the architecture elements, along with its technical layers, will be given in the 2nd iteration of the deliverable.

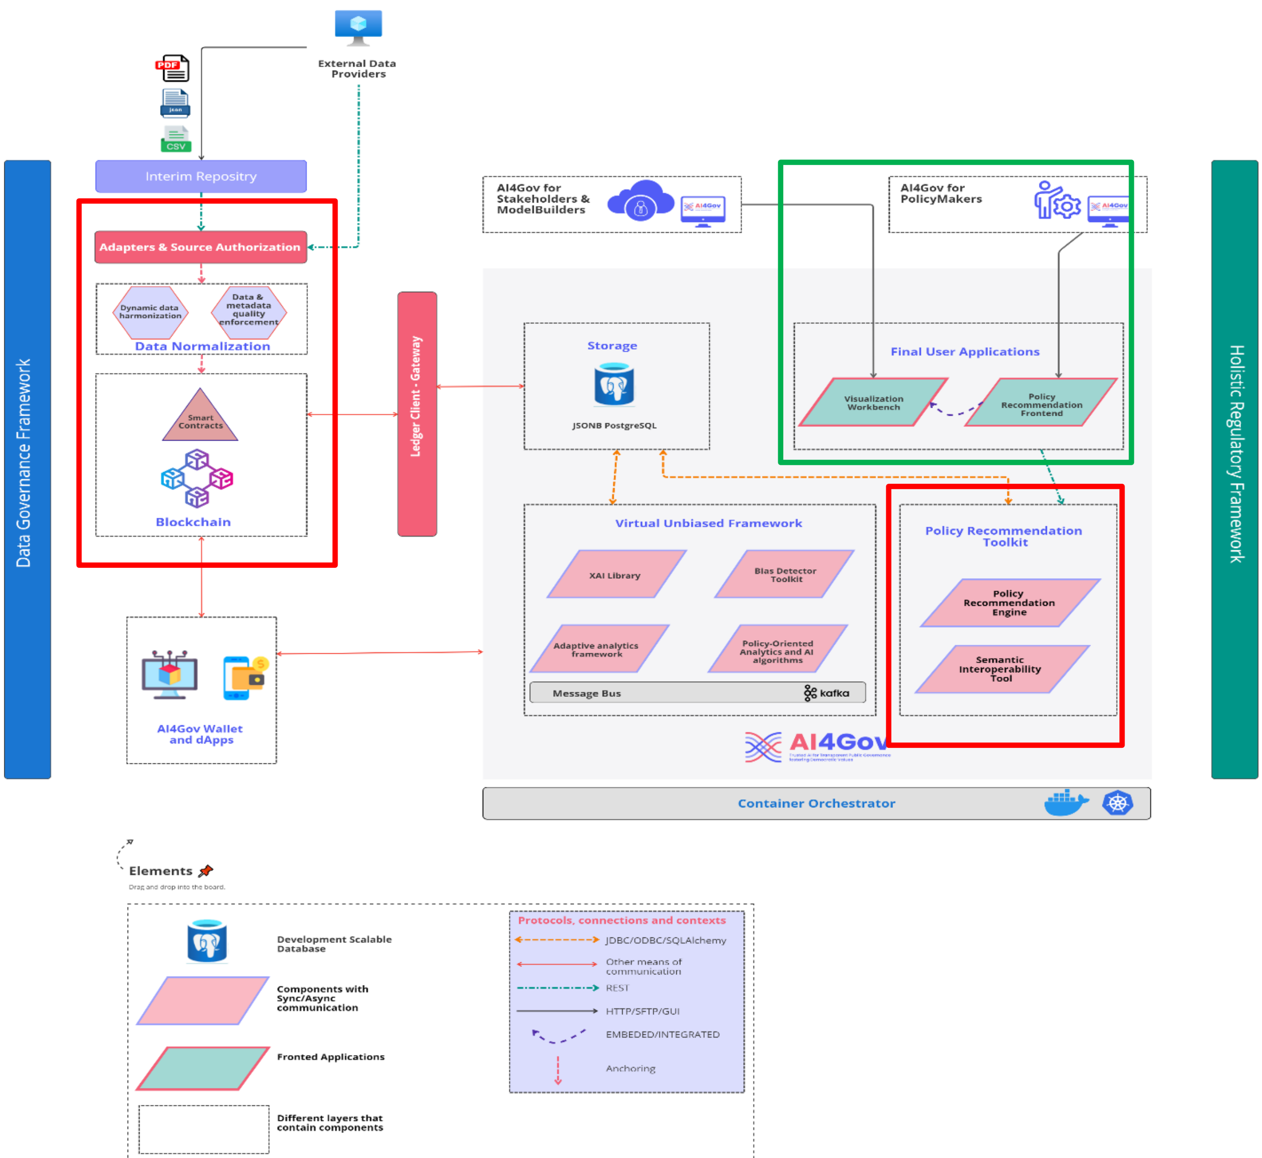

Figure 4: AI4Gov Reference Architecture

3.1 Business Layer

From a business perspective, the main value of the PRT is to facilitate policy making. However, as it was already seen by the requirement analysis, it does so by incorporating various interconnected functionalities and affects the various stakeholders in a different way. To this end, a high-level business viewpoint will be given, that will be accompanied by the different views that describes how these different aspects.

The high level viewpoint is presented in Figure 5. A circular value stream provides clear and unambiguously defined policies that are enhanced by the engagement of the stakeholders. The stakeholders co-create the policies by using a platform that promotes trust. The policies are published and are then governed and audited via the open platform. The business services for materializing this value stream are:

- The Semantic Alignment service is leveraged to produce policies with clear semantics.
- The Open Democracy DAO implements all functionality that allows stakeholders to participate and govern the policy-making process in a democratic and decentralized manner using the underlying blockchain infrastructure.
- The AI Recommendation service uses AI to recommend policies based on criteria set either by the policy makers or by the democratic process agreed upon in the DAO service.

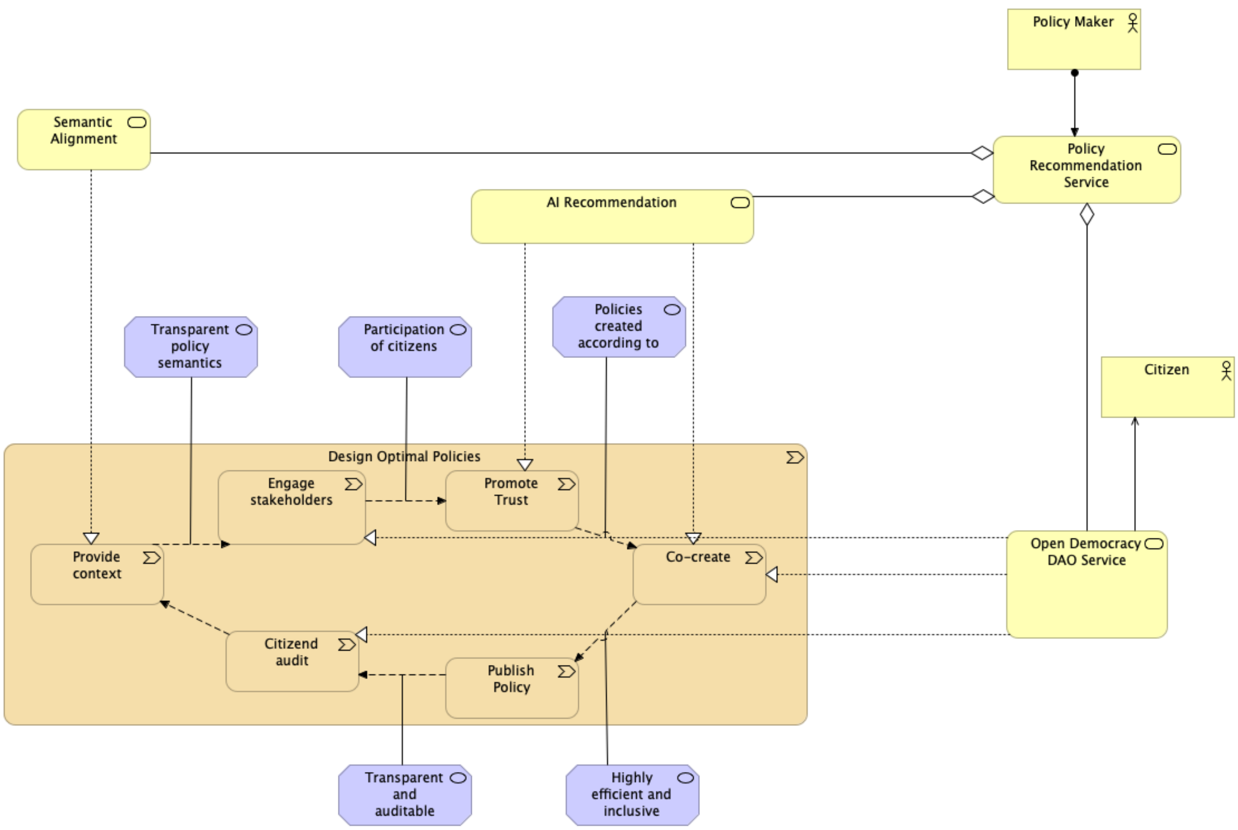

Figure 5: PRT Architecture – High-level view

3.1.1 Semantic Alignment View

Figure 6 depicts the Semantic Alignment View that shows how the Semantic Alignment service is decomposed. It is realized by an underlying business process, which is served by two sub-services. The Data Uplift Service performs the translation of data headers, relations, and metadata from the source format to the common vocabulary, while the Data Alignment Service maps data between data sets using common ontologies.

These services are in turn realized by internal business processes, which are composed respectively by internal business functions which define that uplift and alignment should be performed by using discovery services that are based on AI.

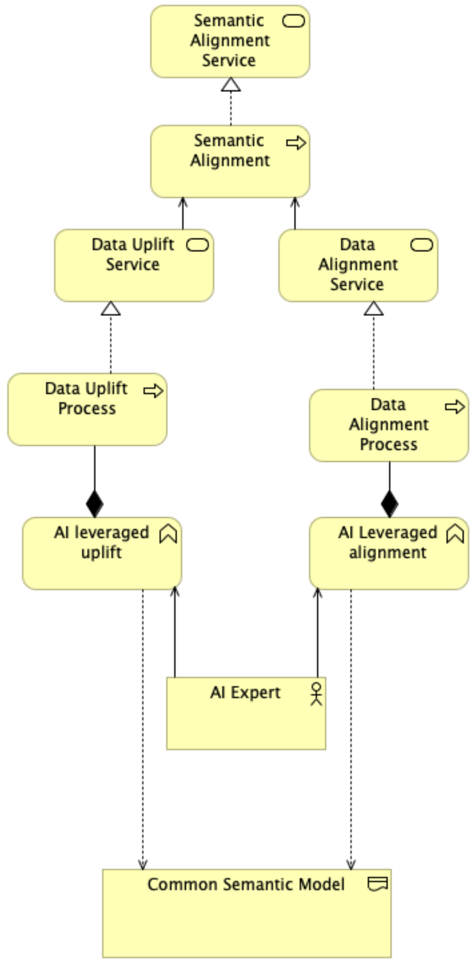

Figure 6: Semantic Alignment View

3.1.2 Open Democracy View

The Open Democracy DAO is a Decentralized Autonomous Organization[1] realized by the elements depicted in the Open Democracy View in Figure 7. The Service is realized by three processes that cover all main aspects of the service, mainly voting, consulting (opinion forming) and auditing. The functions that implement this view are all based on business logic implemented in smart contracts using the underlying Decentralized Data Governance infrastructure.

[1] https://www.investopedia.com/tech/what-dao/

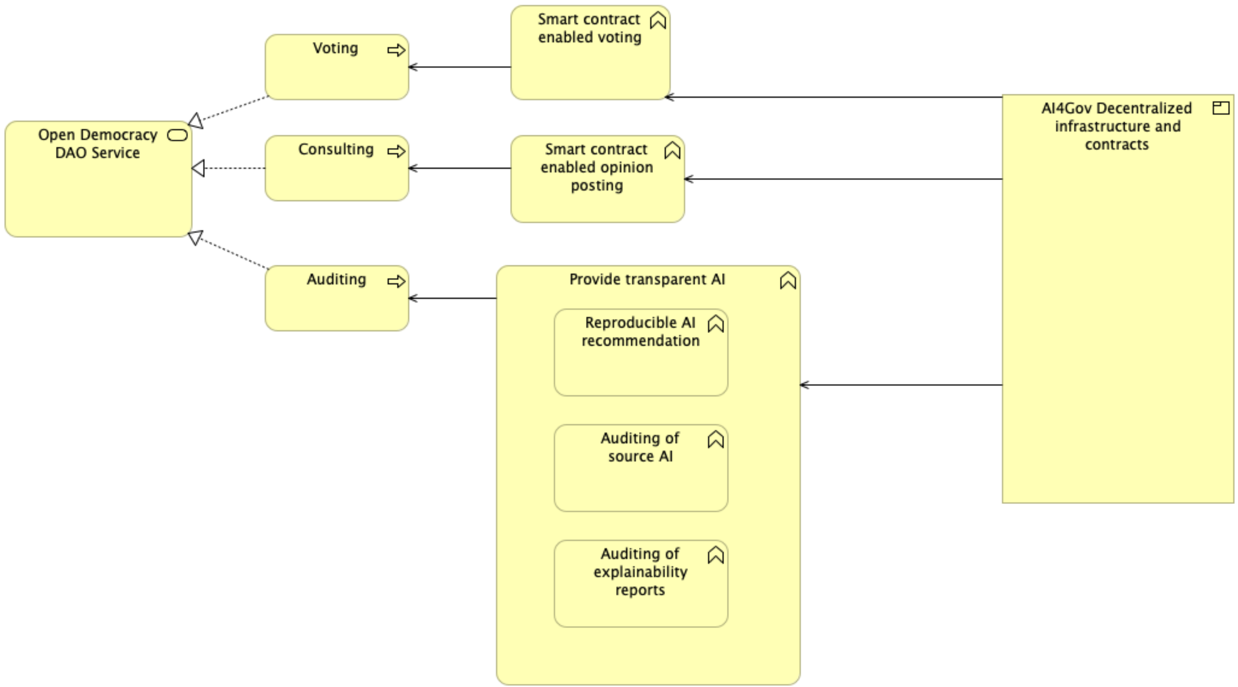

Figure 7: Open Democracy View

3.1.3 AI-based Policy Recommendation

The AI-based Policy Recommendation View is depicted in Figure 8. The realization of the AI Recommendation Service takes place via two functions, mainly the one that is responsible for the recommendation itself and the one that is responsible for anchoring any AI and explaining ability reports that happened off-chain to the policy that was recommended based on this AI. The on-chain AI function is served by the AI4Gov Decentralized Infrastructure and Contracts product (defined in D3.1). External AI Services that are to be implemented in WP4 can be anchored in the blockchain following the processes of data anchoring defined in the Decentralized Data Governance Framework and, in turn, served to the AI Recommendation Service via the On-chain AI function.

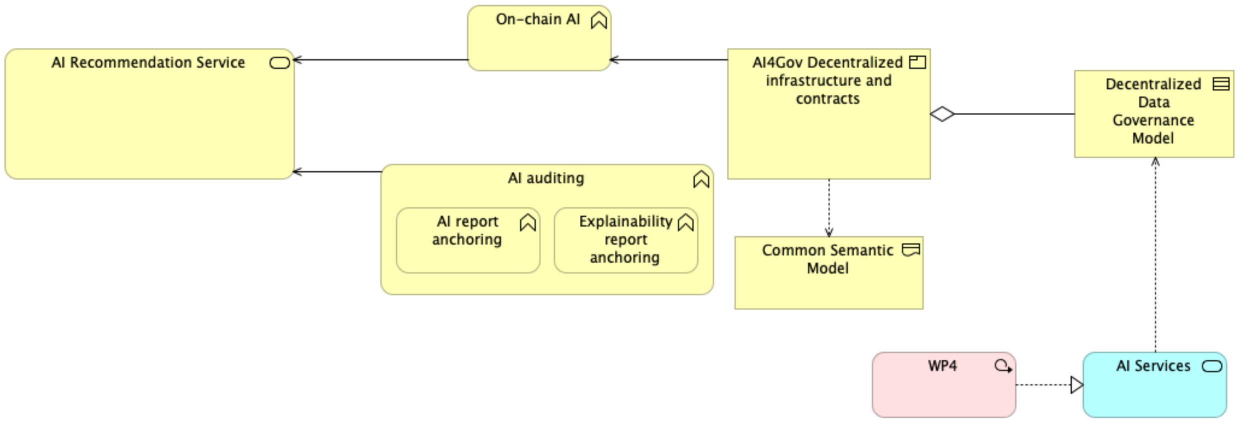

Figure 8: AI Recommendation Service

Although initially the main goal was to link the AI with its output for auditing purposes, the opportunity to implement AI directly on-chain has been identified during the requirements gathering phase. This idea builds upon existing trends of using large blockchain networks to scale AI. Apart from the processing power that a large network can provide, on-chain AI has the added benefit that it can be directly validated by peers by invoking the underlying smart contract(s). LLMs have already been implemented in the Ethereum blockchain[2], and we intend to investigate the idea of developing such LLMs in the HyperLedger Fabric architecture as well. Although the network is small, we anticipate that as large-scale blockchain networks, such as EBSI, are progressively adopted in the EU, such use cases may provide measurable benefits in the future.

3.2 Semantic Layer

In order not to be confused with the Semantic Alignment View, which was described in 3.1.1, the Semantic Layer depicts the organization of data and information, whereas the Semantic Alignment View describes the process of performing semantic alignment. This layer is depicted in Figure 9. Sources of data include:

1. AI and explainability reports that are generated by external AI services and are linked into the blockchain via anchors. The anchoring business function describes this process and is used to serve the recommendation system based on AI.
2. The source and the aligned data come from data that are retrieved on-site; the first one corresponds to raw data as these are collected by pilots, while the latter one corresponds to data that have been uplifted by the semantic models.
3. Finally, the policy is a serialized document that describes policies and is stored on-chain.

[2] https://medium.com/@ModulusLabs/chapter-14-the-worlds-1st-on-chain-llm-7e389189f85e

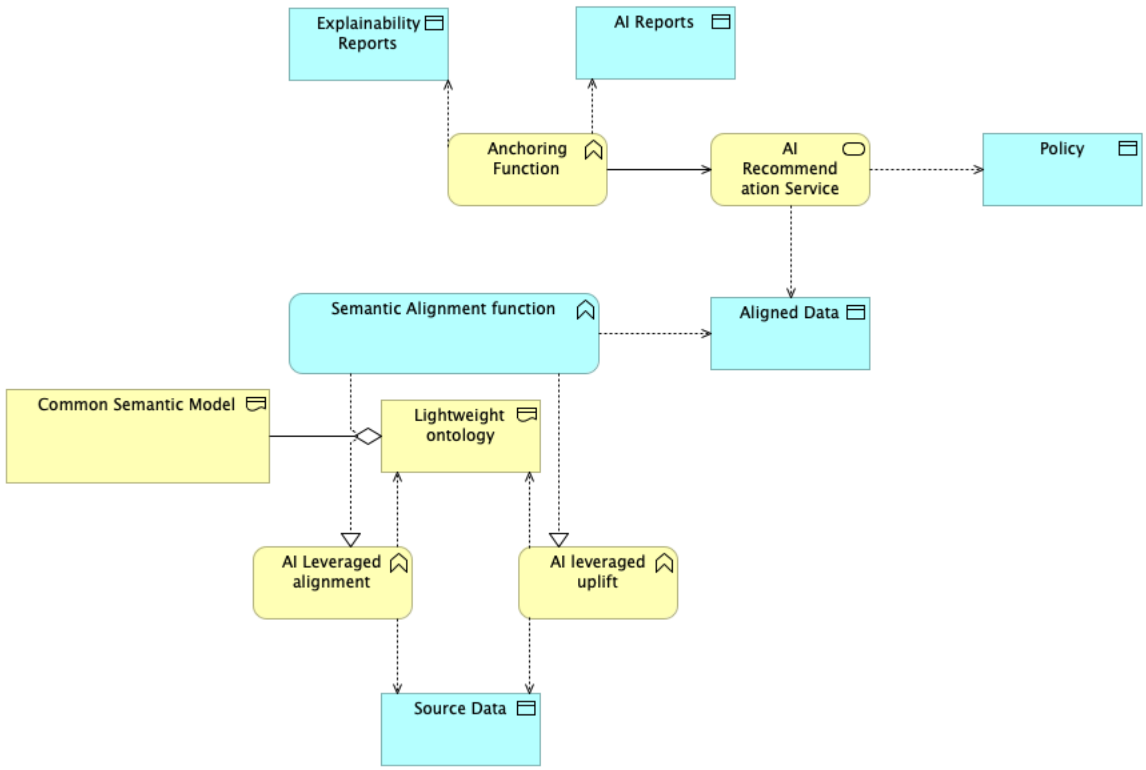

Figure 9: Semantic Layer View

3.3 Application Layer

The high-level application layer of the PRT is depicted in Figure 10. The toolkit is a component that is assigned to two services: the DAO service, which implements the Open Democracy functionality, and the AI Recommendation Service, which implements On Chain AI and mapping of policies to AI reports via blockchain anchoring. For completeness, all relevant business services that are realized by the corresponding application services are also listed in the diagram. The Semantic alignment is not directly assigned to the component but rather serves other functionalities of the PRT, namely the recommendation service itself. The role of the main application services is to serve the core business functions that compose the PRT, such as the auditing/voting systems and the recommender system.

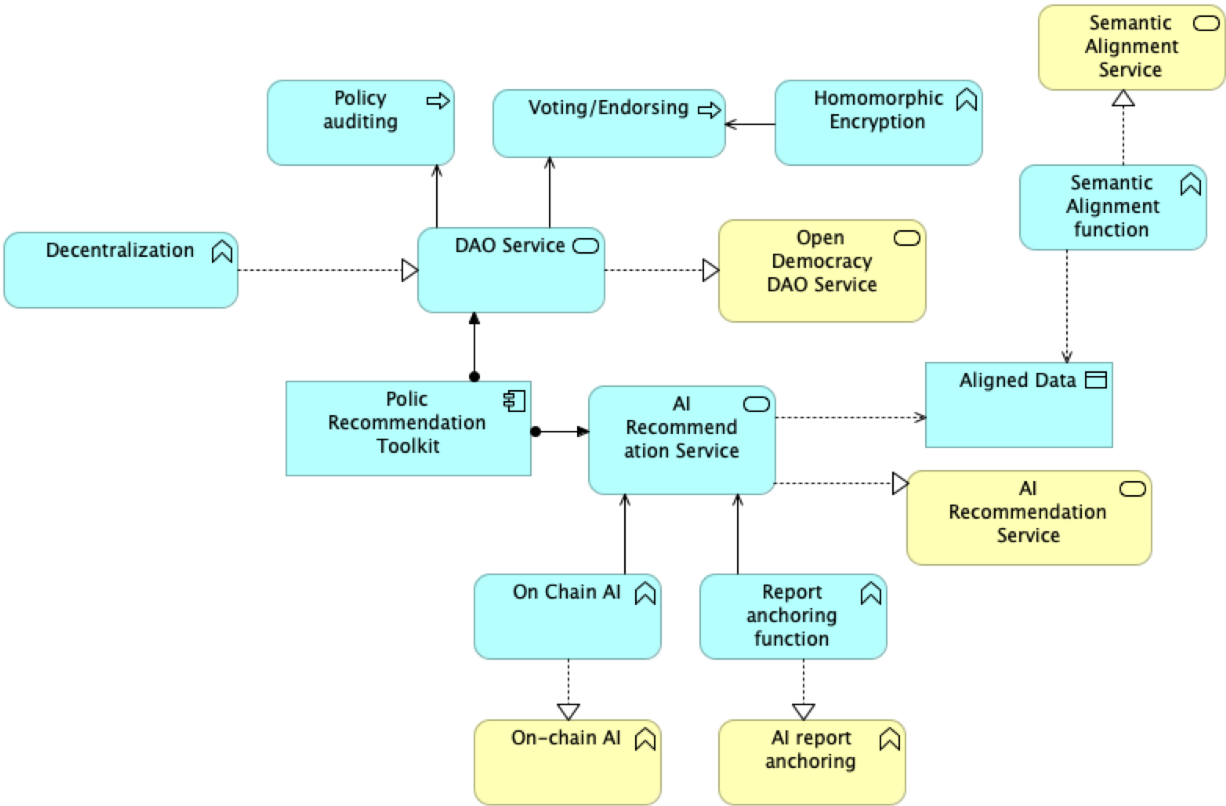

Figure 10: PRT Architecture – Application Layer

4 Prototype Development

The implementation of the PRT is following an iterative process by which a continuous requirements engineering process is followed in parallel with pilot activities (WP 6) to constantly update the requirements and the software releases following suit to implement a subset of these requirements in each cycle.

In parallel with the implementation of PRT, integration activities have also started; both the decentralized infrastructure and the PRT code base are integrated into the Visualization Workbench as far as the desktop elements are concerned. Citizens are equipped with a mobile application implemented as a decentralized application (dApp), which borrows design elements from the Visualization Workbench in order to create a unified user experience.

Regarding the overall status of the implementation in relation to all the functionalities that the PRT promises, this can be summarized as follows:

- The mechanism for aligning policy action items into a lightweight semantic model has been implemented. The model was developed based on policy items from the VVV pilot and is used for prototyping purposes.
- The mechanism for aligning KPIs into a lightweight semantic model has been implemented. The model was developed based on KPIs from the VVV pilot and is used for prototyping purposes.
- Insertion and updating of policies have been implemented as a smart contract using the underlying decentralized infrastructure that is described in D3.1.
- Mechanisms for voting policies have also been implemented. Voting can be performed by both institutions and citizens via their dApps.
- Smart contracts for automatically filtering out popular policies have been implemented. These smart contracts identify popular policies that are endorsed by peers (e.g., other municipalities) or by peers upon conditions set into the smart contract.
- An on-chain recommendation system has been implemented as a smart contract as a prototype of on-chain AI algorithms. The prototypical recommendation system accepts a set of hard and soft KPIs and returns all policies that fulfill the hard KPIs set by the user. The policies are then ranked according to their performance on the soft KPIs.
- An Android application, tested on an emulator that allows citizens to monitor policies and participate in voting and giving feedback on them.

Some representative screenshots of the above functionalities are illustrated below (Fig 11-15). At first, the public servant will log in to the Visualization Workbench and access the toolkit via the appropriate menus. She/he can then initialize a policy by giving a name (Figure 11). The policy can then be given a set of action items along with estimated performance (Figure 12). All data regarding policy will be stored in the blockchain. The recommender system, implemented as a smart contract, can retrieve policies that match a set of hard and soft KPIs (Figure 9); the smart contract definition can be altered to define new rules for filtering and scoring. Finally, citizens or other public institutions can inspect and endorse a policy (Figure 14); policies that gather enough votes can gain special status.

Monitoring and evaluation of policies can also be performed by citizens via an appropriate dApp (Figure 15).

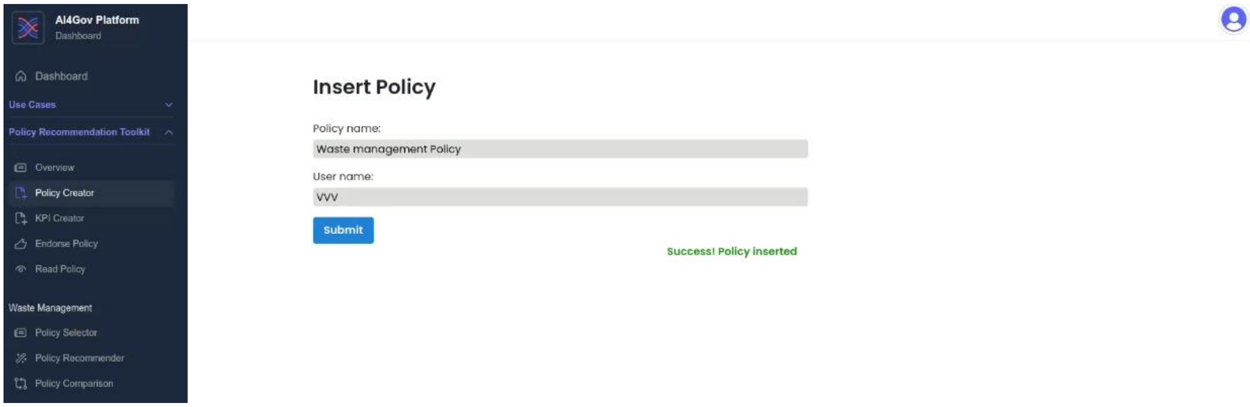

Figure 11: Create a policy and assign a name; the policy is stored in the blockchain

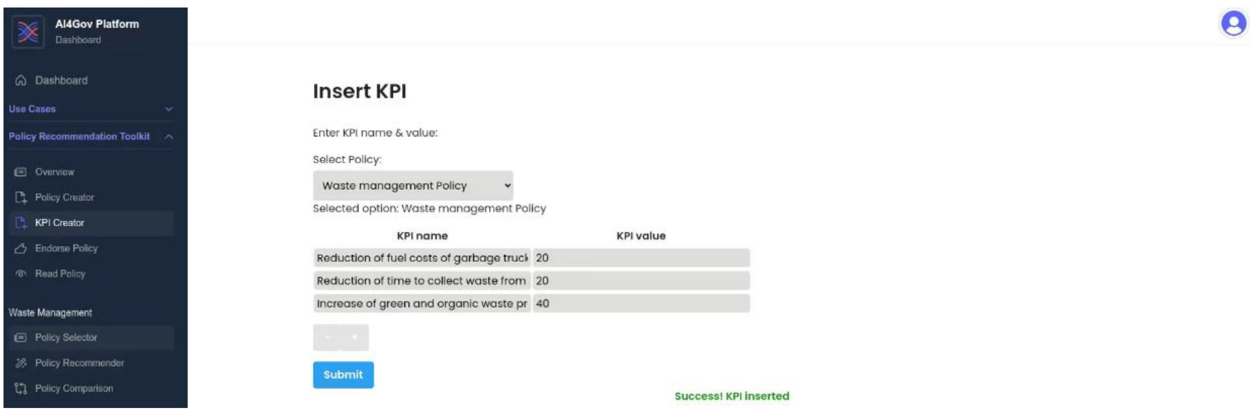

Figure 12: Enter a list of policy action items and the relevant goals using the predefined lightweight ontology. Set values that each item achieves under the policy.

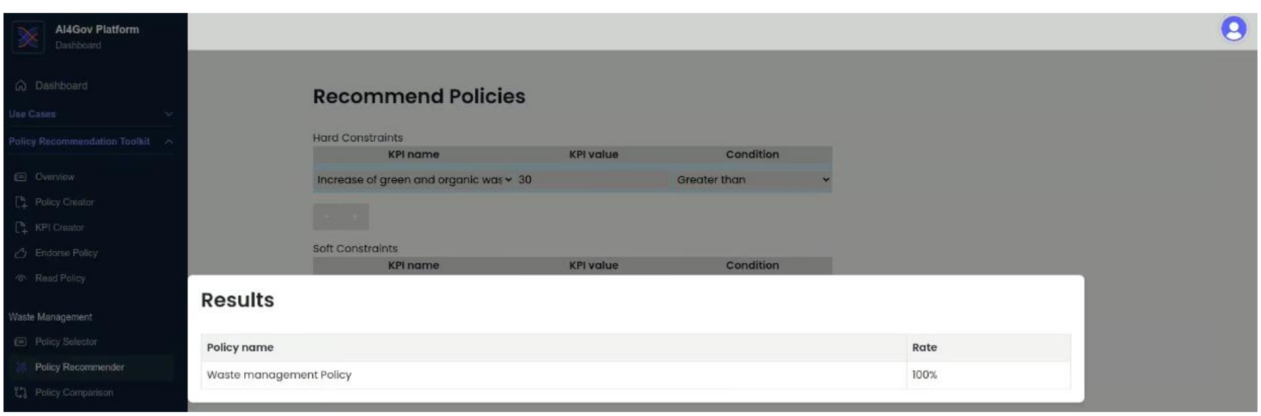

Figure 13: Set a set of hard and soft criteria, and return the policies that match the criteria together with their score. The business logic is implemented via a smart contract.

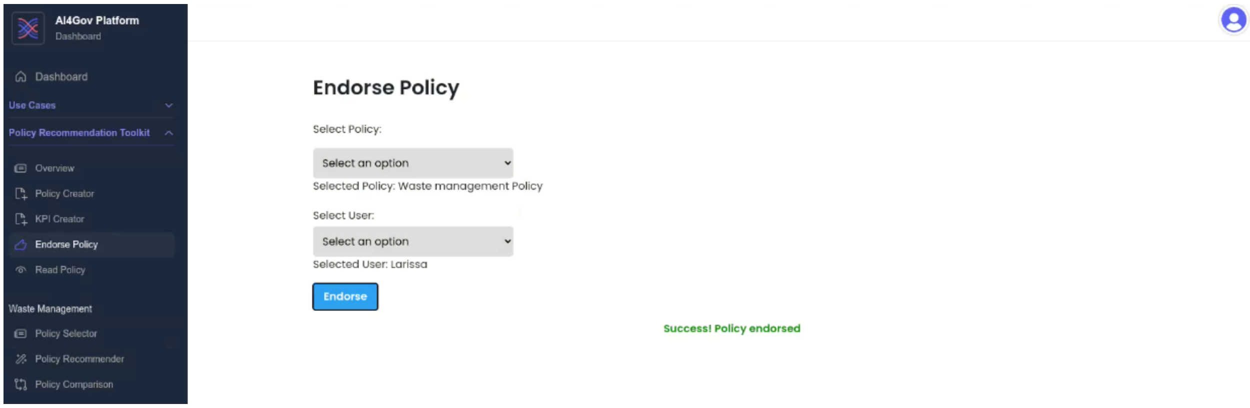

Figure 14: Endorse a policy. If enough votes go to a policy, it will gain a "popular" status

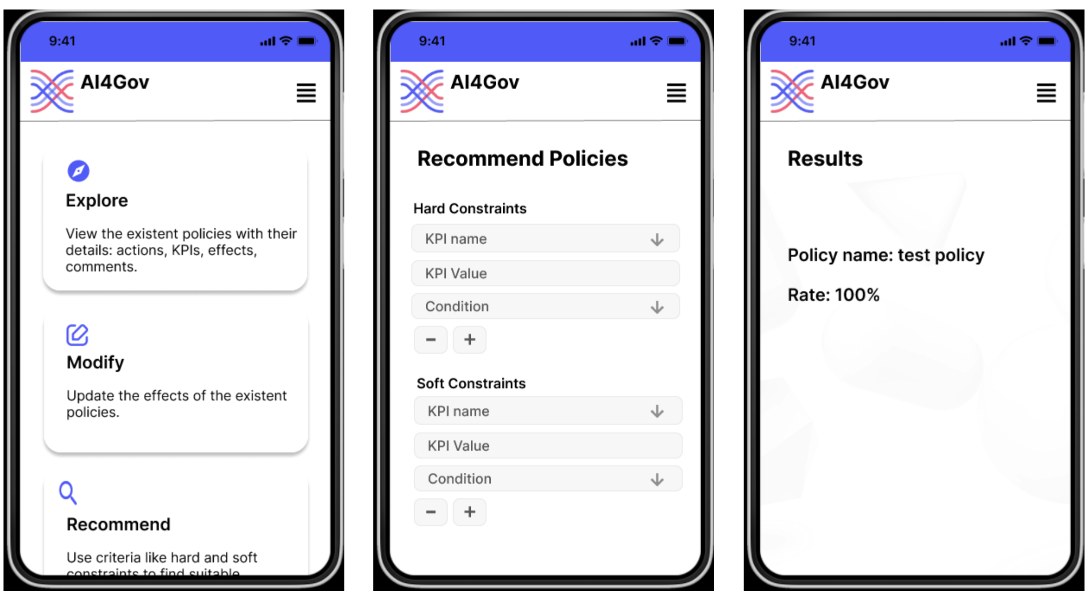

Figure 15: Citizen dApp for viewing and evaluating policies.

4.1 Future work

Future work consists of enriching all the features of the PRT to provide a fully customizable, AI-enabled toolkit for open and inclusive policy-making involving public authorities and citizens.

While the requirement engineering process follows an agile approach and, thus, the requirements are not fully finalized, some core functionalities that need to be implemented can already be identified. These are:

- The inclusion of taxonomies and vocabularies into the Decentralized Governance Model and the utilization of those taxonomies to define action items and KPIs of policies.
- The implementation of an AI-enabled schema aligner will map entities from source data to the taxonomies and vocabulary of the PRT.
- The implementation of mechanisms for defining custom blockchain governance models and deploying DAOs on the blockchain. With these mechanisms, public authorities can initiate consultation processes and define rules of participation and rules of voting.
- The implementation of a mechanism for automated smart contract definition and deployment. This task is common with activities carried out in T3.1; in the context of T3.3 it will focus on applications that allow users to set the basic rules and meta-parameters of a recommender smart contract (e.g., threshold values for accepting policies, sequence of conditions etc.) and govern how this updated business logic is deployed in the blockchain.
- The investigation of the deployment of complex on-chain AI algorithms that will offer transparent and reproducible results. Specifically, we investigate the possibility of deploying on-chain LLMs in the HyperLedger Fabric infrastructure.
- The implementation of fully autonomous citizen dApps that will be packaged and deployed on Android devices. As was explained in D3.1, we intend to make the citizen wallet that will this and any other dApp EBSI-conformant.

5 Conclusions

In this report, the first iteration of the requirements of the Policy Recommendation Toolkit was given, together with the first version of the architecture that fulfills these requirements. As efficient policymaking should allow for citizen feedback and co-creation, the requirements defined in the initial pilot definitions have been expanded, as to include a citizen component that is expected to actively increase citizen participation in policymaking. Although many of the technology enablers that promote openness and inclusiveness have already been documented in D3.1, new enablers have been identified that are exclusive to the use cases of the PRT, namely on-chain AI, Zero Knowledge Proofs, and Homomorphic Encryption. These enablers further promote transparency and trust by allowing on-chain execution and validation of AI algorithms and securing the secrecy of votes of citizens when this is demanded.

6 References

Feige, U., Fiat, A., & Shamir, A. (1987). ZERO KNOWLEDGE PROOFS OF IDENTITY. *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, 210–217. https://doi.org/10.1145/28395.28419

Landemore, H. (2020). *Open democracy: Reinventing popular rule for the twenty-first century*. https://books.google.com/books?hl=en&lr=&id=fe7xDwAAQBAJ&oi=fnd&pg=PA247&dq=open+democrcy+helene+landemore&ots=sw3kkUpnpM&sig=vdlESk8SoBWBN9XU_zEM9PuJUt0

Semantic Web - Wikipedia. (n.d.). Retrieved April 22, 2024, from https://www.w3.org/standards/semanticweb/

Thatcher, M., & J. B.-I. E. of the S., & 2015, undefined. (n.d.). Issue Networks: Iron Triangles, Subgovernments, Policy Communities, Policy Networks. *Iris.Luiss.ItM Thatcher, J BraunsteinInternational Encyclopedia of the Social & Behavioral Sciences (Second Edition), 2015•iris.Luiss.It*. Retrieved April 19, 2024, from https://iris.luiss.it/handle/11385/184949