



Deliverable 3.2 Decentralized Data Governance, Provenance and Reliability V2


30-06-2024

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	Ubitech
License	 <p>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/</p>

AUTHORS		
	Name	Organisation
Document leader	Ntalaperas Dimitris	Ubitech
Participants	Xanthi Papageorgiou	Ubitech
	Nikos Kalatzis	Ubitech
	Konstantinos Tzelaptops	Ubitech
	Dimitris Kotios	UPRC
	George Manias	UPRC
Reviewers	Matej Kovačič	JSI
	Sotiris Athanassopoulos	MAG

VERSION HISTORY				
Version	Date	Author	Organisation	Description
0.1	18 May 2024	Ntalaperas Dimitrios, Kotios Dimitris	UBI, UPRC	Table of Contents
0.5	23 May 2024	Ntalaperas Dimitrios, Xanthi Papageorgiou	UBI	SotA and Architecture
0.7	1 Jun 2024	Ntalaperas Dimitrios, Kalantzis Nikos, Tzelaptsis Konstantinos	UBI	Prototypes
0.8	18 Jun 2024	Kotios Dimitris	UPRC	Data Governance Framework
0.9	22 Jun 2024	Ntalaperas Dimitrios	UBI	Editing and first draft for internal review
1.0	30 Jun 2024	Ntalaperas Dimitrios	UBI	2 nd draft incorporating internal review comments

Table of Contents

Abstract	6
1 Introduction.....	7
1.1 Purpose and scope.....	7
1.2 Document structure.....	7
1.3 Updates with respect to previous version	8
2 State of the Art	9
2.1 Decentralized Autonomous Organizations (DAOs)	9
2.1.1 <i>The DAO</i>	13
2.1.2 <i>The Terra LUNA disaster</i>	13
2.1.3 <i>Protecting citizens in a DAO</i>	15
2.2 Markets in Crypto-Assets (MiCA) Regulation	16
2.3 HyperLedger Aries.....	17
2.4 Transparency of AI	17
3 Decentralisation in AI4Gov	19
3.1 Data information in AI4Gov	19
3.2 Decentralised Data Storage in the AI4Gov platform	21
3.3 Architecture for Decentralized Data Governance	23
3.3.1 <i>Business layer</i>	23
3.3.2 <i>Application layer</i>	26
3.3.3 <i>ABB viewpoint</i>	26
3.3.4 <i>SBB viewpoint</i>	28
3.3.5 <i>General requirements</i>	33
4 Technological enablers	35
4.1 The HyperLedger Aries.....	35
4.2 Prototypical implementation	36
4.2.1 <i>Decentralized Policy Making</i>	37
4.2.2 <i>Identity Management and Verifiable Credentials</i>	39
5 Data Governance Framework	44
5.1 General Guidelines and Policies.....	44
5.2 Applicable Regulations and EU Guidelines.....	47
5.2.1 <i>General Data Protection Regulation (GDPR)</i>	47
5.2.2 <i>EBSI Conformance</i>	50
5.2.3 <i>Ethics Guidelines for Trustworthy AI</i>	50
5.2.4 <i>EU Artificial Intelligence Act</i>	52
6 Conclusions.....	55
7 References	56

List of figures

Figure 1: Blockchain enabled governance for public organizations	11
Figure 2: Fully functional DAO for determining policies. A member can propose a change, if this is approved by the governing body, it is then sent back to members for voting.....	12
Figure 3: AI4Gov Reference Architecture	22
Figure 4: 1 st iteration of the Business layer of the decentralised architecture	24
Figure 5: Motivational viewpoint for the decentralized architecture	25
Figure 6: Business viewpoint of the decentralized infrastructure.....	26
Figure 7: 1 st version of the Application layer of the decentralised architecture	26
Figure 8: ABB viewpoint of the decentralized architecture.....	28
Figure 9: SBB diagram for the Decentralization function and the Identity Management and VCs process.	29
Figure 10: SBB diagram policy administration functionality.....	31
Figure 11: SBB diagram for content assessment	32
Figure 12:SBB for the technical infrastructure viewpoint	33
Figure 13: The HyperLedger Aries stack	36
Figure 14: PRT – Insert Policy Screen.	37
Figure 15: Policy created successfully.	38
Figure 16: Association of KPIs to a policy	38
Figure 17: Defining criteria for policy recommendations.....	39
Figure 18: PRT recommendations using a smart contract.....	39
Figure 19: Citizens’ wallet – Initial screen	40
Figure 20: Boarding invitation generated by HyperLedger Aries	41
Figure 21: Accepting the invitation	42
Figure 22: Accepting a credential offer	43
Figure 23: AI Act defined levels of risk	52

List of Tables

Table 1: Types of data and end users per pilot case.....	19
Table 2: Decentralised Data Governance policies in AI4Gov	34

Abbreviations

Abbreviation	Description
ABB	Architecture Building Block
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
IPFS	InterPlanetary File System
PoS	Proof of Stake
PoW	Proof of Work
SBB	Solution Building Block
SSI	Self-Sovereign Identity

Abstract

The present document presents the second iteration of the Decentralised Data Governance (DDG) Model for AI4Gov and describes the mechanisms for achieving and implementing the required provenance and reliability features of AI4Gov while also respecting the privacy guidelines set by GDPR. The second iteration of the DDG model involves two major items. Firstly, the finalization of the architecture and prototypes that allow fully transparent data governance together with the definition and execution of custom decentralized business scenarios using smart contracts. Secondly, and most importantly, there was a deliberate redesign and refactoring on both the codebase and the architecture to accommodate a citizen-centric approach that will allow citizens to board the platform and participate in collaborative and co-creative processes under a framework that supports procedures that are core to open democracy. This new direction demanded some items of our first approach, which was documented in D3.1, to be reconsidered. These were:

- Redesign the architecture and implementation to accommodate citizen wallets. While the OpenDSU is a good candidate solution for organization wallets, it was proven to create wallets that required resources, making the citizen dapps slow. To this end, the implementation leveraged frameworks such as HyperLedger Ares instead, which are ideal for decentralized identity management, especially when data subjects are not known beforehand, as is the case with citizens boarding a decentralized infrastructure.
- The implementation of fully self-governed units, in which the business terms, the governance model, and the rules by which this could be changed are clearly defined upon creation. These units follow the Digital Autonomous Organization (DAO) paradigm that is already followed in many public blockchains.

The current report will review the state-of-the-art areas that concern DAOs and will also make a short report on the new MiCA regulation and the ways this new regulation framework could potentially affect any open solution democracy. The new revised architecture will be described, with an emphasis on the changes needed to include the new functionalities for supporting open democracy. The new technology stack incorporating the Aries framework is also presented, demonstrating how citizens can be verified in the blockchain using their credentials and, inversely, how the blockchain can provide proofs to citizens that the citizens can verify.

1 Introduction

1.1 Purpose and scope

The present document offers the 2nd iteration of the decentralised data governance framework of AI4Gov. Its major constituents are the final specifications, which, build upon an expanded description of use cases that also leverage open democracy. From this description an expanded set of technical requirements is derived that naturally lead to a new architecture.

The State-of-the-Art analysis that was performed in D3.1 remains valid in the 2nd iteration; however, two extra sections are added in this iteration that are relevant to principles of citizen-centric and citizen-governed blockchains. The first one is centered around Digital Autonomous Organizations (DAOs) and explains both the operating principles and the dangers that such platforms can present, together with common mitigation measures. The second one concerns the new MiCA regulation; although it is a regulation about crypto-currencies, it can be relevant to governance models that are based on tokens. Therefore, we briefly review MiCA together with potential applications in EU-based eGov solutions based on blockchain. Lastly, the state-of-the-art synergistic solutions containing both AI and blockchain are also briefly overviewed; the way that blockchain can leverage existing AI techniques to further promote transparency in AI is briefly described.

The final version of the architecture detailing the technical layer is presented with a prototype that uses HyperLedger Aries and HyperLedger Fabric to cover both organizational and citizen needs. Briefly, Ares is used for identity and wallet management at the user level and Fabric at the organizational level. The synergy of the two frameworks and the way it fulfills the old and new requirements are presented in the architecture and is demonstrated via two demonstrator scenarios.

Finally, the updated and last version of the Data Governance Framework (DGF) is presented.

1.2 Document structure

The present document is structured as follows:

- Section 1 contains the present introduction
- Section 2 gives a State-of-the-art (SotA) analysis of the DAO mechanism, which is central to the implementation of citizen-centric decentralized solutions. Aspects of the new MiCA regulation that may affect the solution in the future together with current approaches of leveraging AI using blockchain are also described.
- Section 3 applies the new disciplines analyzed in Section 2 and amends the previous data policies accordingly. It presents the new version of the architecture that incorporates DAOs and provides a description of the application components, which are now instantiated as Solution Building Blocks (SBBs).

- Section 4 describes HyperLedger Aries, the new technology enabler used for identity management and Verifiable Credentials. It also presents two representative scenarios that demonstrate all functionality implemented in the framework thus far: the first scenario involves policy-making and recommendations between organizations, and the second one involves citizen participation and credential presentation.
- Section 5 describes the second iteration of the Data Governance Framework of AI4Gov, i.e., the framework that entails the processes and definition that govern all data handled in the project.
- Section 6 gives the conclusions of the present work.

1.3 Updates with respect to previous version

This is the second iteration of the Decentralized Data Governance Model. It provides new state-of-the-art involving DAOs, the MiCA regulation and synergistic solutions involving blockchain and AI. The architecture is finalized and is expanded to incorporate DAO-based solutions for citizen participation. The new Hyperledger Aries technological enabler for Identity Management is also documented. Moreover, the final version of the Data Governance Framework is presented. Eventually, the final prototypes for the decentralized data governance framework are presented.

2 State of the Art

This section is intended to be part of a single State-of-the-Art section that covers all aspects of Blockchain technology, especially those relevant to eGov and transparency in data governance. The basics of blockchain technology, the various types of blockchains (public vs private), the mechanisms for anchoring files into the blockchain, and the execution and governance of smart contracts are covered in Section 2 of D3.1, which is the first iteration of the present work. Here, we will focus mainly on DAOs and the MiCA regulation. At the technology level, and since our new solution leverages HyperLedger Aries¹ for identity management, a review of this technology will also take place.

2.1 Decentralized Autonomous Organizations (DAOs)

One of the original visions behind the creation of many cryptocurrencies, especially Bitcoin and Ethereum, was to provide a means for currency owners to perform transactions that were perfectly decentralized. This involves a shared ledger and mechanisms for reaching a consensus regarding the transactions, as we also analyzed in D3.1, but for a full economy, it goes beyond that. In a capitalistic economy, regardless of the level of state control the economy allows, major decisions regarding macroeconomics are deferred to governance agencies. Most importantly, the amount of currency circulated is based upon decisions and actions of the Central Bank². The currency owners do not have a direct influence on how or how much money is created³. Other aspects of transactions, such as VAT, are similarly outside the influence of the currency holders.

How are these aspects handled in a completely decentralized economy? Clearly, there should be a mechanism for creating a currency that is not inflationary or deflationary. Considering that the dynamics of the blockchain economy can change, this mechanism should also be adaptive. Likewise, there may be a need for a *taxing* mechanism in transactions both for inflation control and for locating cryptocurrency rewards to support whatever reward system the blockchain protocol implements.

One mechanism is to implement this directly into the blockchain protocol, hoping that the rules are generic enough to always achieve the perfect balance. The Bitcoin protocol is such an example. We will not go into details on how Bitcoin maintains asset supply, but we will mention that it clearly defines how and when bitcoins are created and how they are rewarded; the

¹ <https://www.hyperledger.org/projects/aries>

² The government, depending on the legislation, can legally or illegally influence these decisions. For the purposes of this discussion this is irrelevant; the money supply is still governed by mechanisms outside the hands of the currency owners.

³ Here we explicitly refer to direct means of influence. Citizens, organizations and parties can of course vote and lobby in a democracy.

mechanisms are built into the protocol so any participant, either trader or miner, is fully aware of them and has guarantees that these rules are executed exactly as they are written⁴.

What happens if, for whatever reason, reality surpasses these rules and the mechanism is deemed to be insufficient by the community? Can it change on-chain? The answer is no. As we have reviewed in D3.1, the only way to migrate to a new set of rules is by forking. That entails these steps:

- A critical part of the community identifies the need for change.
- The change is incorporated into the software that implements the blockchain protocol.
- Members of the community start using the new protocol by creating a fork of the blockchain. As this is being performed, members that adopt the new protocol will also use the new fork.

This transition depends critically upon acceptance from the community. If the majority agrees, the 'new' fork will be dominant, and the 'new' cryptocurrency will be designated with the same name⁵. This mechanism of updating the blockchain protocol is clearly an off-chain procedure. In a sense, it is not completely decentralized since, however democratic it may be, it was conducted off-chain without being verified by a blockchain mechanism such as a smart contract.

It was exactly to handle such scenarios that the DAOs were inspired and created. DAOs are possible⁶ in blockchain protocols that support Turing-complete calculations, such as the Ethereum blockchain. DAOs are organizations that are fully implemented using smart contracts. Depending on the specific DAO, it can define rules that govern all aspects of the transaction, such as fees, conditions that must be met before a transaction takes place, currency supply, etc. What is more, a DAO can implement mechanisms that define how these rules can change based on the actions of the DAO collective. For example, a DAO can implement a mechanism that allows participants to propose some change (e.g., a change in transaction fees) and a mechanism for endorsing or rejecting such proposals (e.g., a majority voting). Such mechanisms have already been implemented in AI4Gov and have been tested in the domain of policymaking. For example, depicts such a scenario in which an organization within a group of organizations that shares a blockchain can propose some new policies, and the other organizations can vote on the new propositions.

It can be argued that such mechanisms, as those already presented in D3.1, constitute a DAO; however, this could be true only by a purely technical definition. The vision behind DAOs is to provide a means of self-governance to end users. From an eGov perspective, DAOs are the means by which citizens can participate in an inclusive and open scheme of governance. This requires a different vision behind the governance model of the blockchain and a different technical solution

⁴ Provided of course, that the blockchain is not compromised via a 51% attack; see D3.1 for a more detailed review of these mechanisms and potential compromises.

⁵ We have described this mechanism in more detail in D3.1. The 'old' fork may still be used by those that disagree with the fork. As it happens with every split in human history, the majority tends to keep the name.

⁶ In theory, DAOs could be implemented in smart contract languages that are not Turing-complete. However this can be extremely difficult and, in any case, the reason a blockchain protocol supports Turing-complete languages is exactly this, to allow for such custom functionality.

to provide the capability for citizens to identify themselves and prove aspects of their identity⁷. Figure 2 shows a scenario more typical in DAOs. A member proposes a change, such as the modification of tax rates. A special body⁸ of the blockchain approves or disapproves the change; if it is approved, it is sent back to the community for voting.

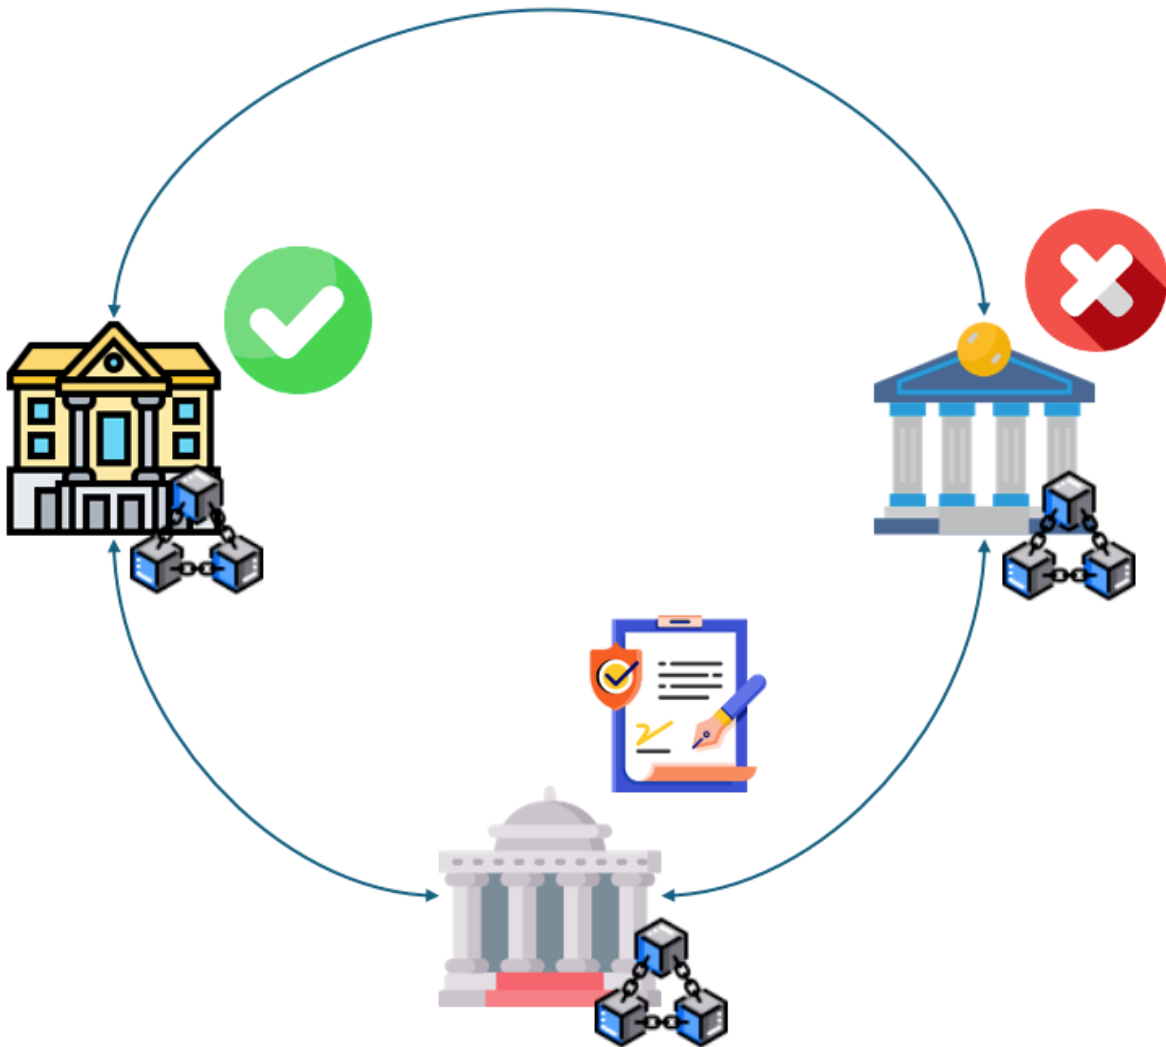


Figure 1: Blockchain enabled governance for public organizations

⁷ It is easy to prove the identity of the organization by providing their certificate and using the MSP mechanism (see D3.1 for an explanation of MSP under HyperLedger Fabric). Allowing for arbitrary citizens to identify and prove their attributes requires a different process.

⁸ Other DAOs may not involve special entities and treat all members equally. The important things that defines a DAO is that rules are transparent and can change by members' actions.

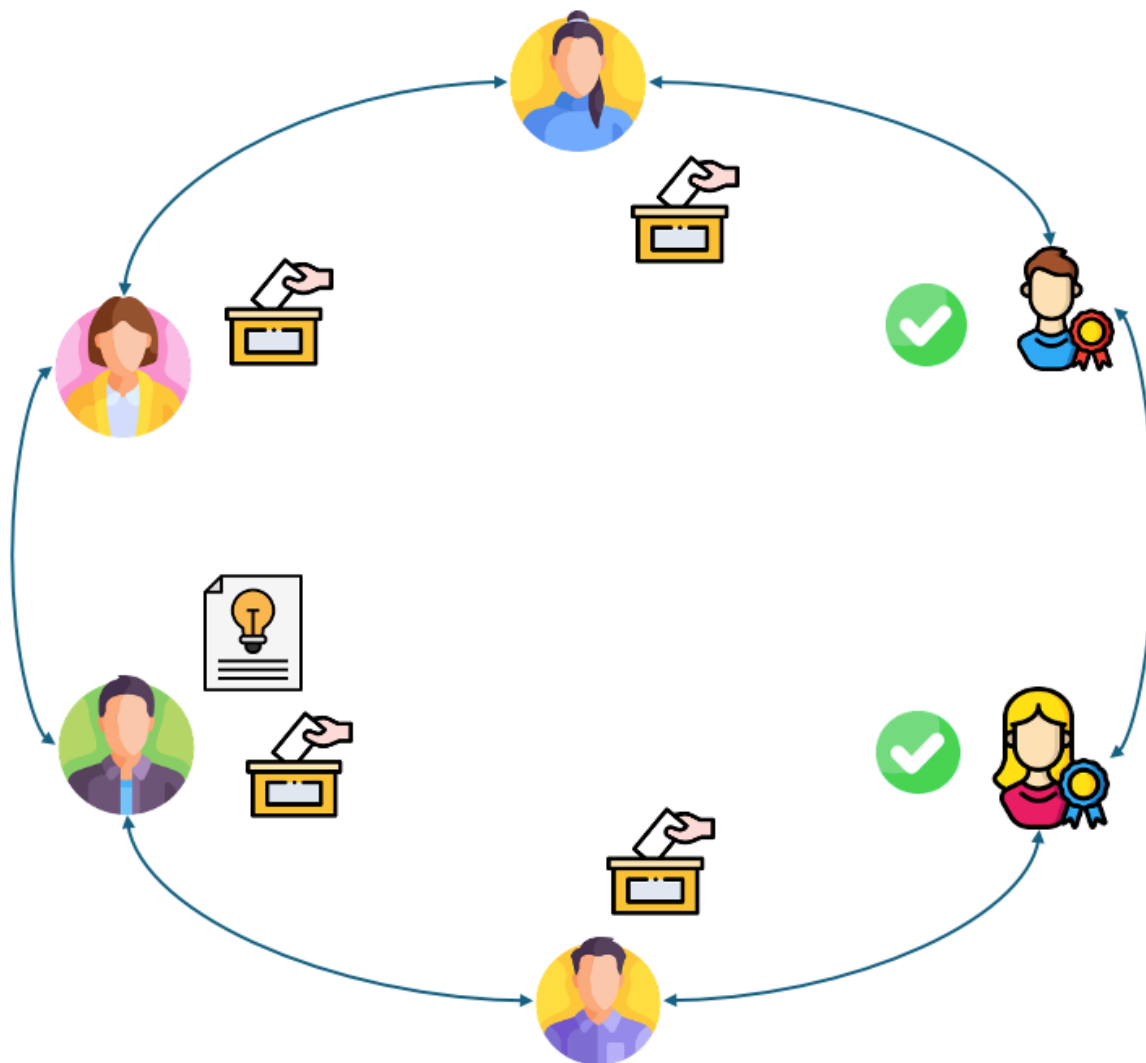


Figure 2: Fully functional DAO for determining policies. A member can propose a change, if this is approved by the governing body, it is then sent back to members for voting.

DAOs seem perfect tools to decentralize an economy⁹. Expanding this into politics, it may seem that DAOs can also be used to decentralize democracy. This may be a misnomer, since democracy is, by definition, a decentralized system of governance. However, aspects of it, such as participation in opinion forming, transparency of processes, verification of candidate claims etc, may be hindered or obfuscated due to various reasons, such as real-life limitations, limited access to original sources, difficult to identify deep fakes etc.

⁹ Whether this is a good idea or not is an area of debate. The fact is that, if the target is decentralization of economy, DAOs are the means to do it.

DAOs can help citizens to somewhat limit the effect of such hindrances by facilitating transparency in decision making and providing more tools for citizen co-creation in policy making. However there are some dangers that need to be mitigated before fully incorporating DAOs in public decision making. These dangers are best described by examples that will be presented in the following two subsections.

2.1.1 The DAO

One of the first DAOs implemented was called *The DAO*, and it was essentially a venture capital fund that was built upon the Ethereum blockchain. It used its own token, called DAO. It concentrated the attention of many investors, and it reached a capitalization of roughly 150 million dollars by 2016. *The DAO*, like all DAOs, was implemented as a smart contract, and its source code was open¹⁰. A number of vulnerabilities in the code led to the development of a hack that transferred around 3.5 Ether coins of the total 11.5 Ether that were invested in the DAO to a third-party wallet; essentially, a third of the total Ether was stolen.

It is to be noted that under the rules specified by the DAO, this was a perfectly valid transaction. Smart contracts and the blockchain cannot differentiate normal from malicious intent. The developers' intentions are irrelevant under this setting. If the DAO's smart contracts allow for a transaction, this transaction is validated by the blockchain's consensus mechanism and is final. Moreover, it cannot be undone by any mechanism within the DAO.

The situation was, in the end, remedied by having the whole Ethereum blockchain forked from a point before the hack took place, thus restoring the 'hacked' wallets¹¹. The hard fork was adopted by the majority of the community. A minority continued to use the old one; the associated cryptocurrency is now called Ethereum Classic. This off-chain mitigation action caused controversy in the Ethereum community. However, it prevented the loss of property that occurred via the hack.

2.1.2 The Terra LUNA disaster

The Terra LUNA disaster is a good example of how a poorly defined DAO can lead to a runaway condition with catastrophic consequences. The Terra LUNA ecosystem was developed by Terraform Labs. It operated on the Terra blockchain network and consisted of two cryptocurrencies:

- The UST, TerraUSD, or just Terra, was a stablecoin, which was pegged to the dollar.
- The LUNA token which was the Terra's blockchain native cryptocurrency used for all the normal blockchain operations, such as paying transaction fees, staking to validate transactions, voting in DAOs, etc.

¹⁰ To anticipate any reasoning that 'closing' the source could have averted the disaster, it should be noted, as was also written in D3.1, that all code in the Ethereum blockchain is in the end deployed as EVM code in blocks for everyone to see. Closing the source can only delay the discovery of vulnerabilities.

¹¹ And also invalidating any transaction from this point on.

Stablecoins are typically pegged to the dollar and have a fixed value of 1 dollar per token. Usually, this peg is maintained by the governing body of the coin by having an amount of off-chain liquid assets that backs its value¹². UST was different in this aspect, in that it retained its peg by a purely algorithmic process.

At the core of this algorithmic process were smart contracts that allowed the exchange of LUNA and UST at the pegged rate of 1 dollar of UST to the amount of LUNA that corresponded to the swapped price of UST. When the UST price exceeded 1 dollar, the users would buy LUNA and convert LUNA to UST with the aim of selling the UST and pocketing the difference. Since the UST traded at above 1 dollar the net difference of total LUNA coins in circulation would decrease. On the contrary the total amount of UST would increase thus creating an inflationary pressure to the coin which would drive its price to fall to the 1-dollar mark.

In the case when UST fell below 1 dollar, the reverse process would take effect; users would buy UST and convert it to LUNA and sell it. Due to the price difference, the total UST supply would decrease, thereby creating deflationary pressure that would drive its price towards the 1-dollar mark.

In early 2022 there was a general decline in cryptocurrency market prices. This drove the price of LUNA down until, at some point, its total capitalization was equal to that of the total UST capitalization. This caused investors to think that the total value and volume of LUNA was no longer enough to support the UST peg¹³, as dollar to dollar there were not enough LUNA tokens to redeem UST tokens.

The Terra LUNA run started with two massive withdrawals of UST from a Terra DAO^{14 15}. This process initiated a drop in the UST price at around 0.90 cents of the dollar. As we have seen, this should have triggered the users to swap UST to LUNA and sell it. These transactions of massive volume did not give the chance to the algorithm to adjust the values. The value of LUNA plummeted, both due to the general cryptocurrency crisis, but also due to the fact that huge amount of LUNA tokens were created by users in order to swap USTs. UST holders also started to sell UST directly, even at a loss, to minimize loss. This led to the joint collapse of UST and LUNA prices.

¹² More precisely, the price is maintained by the amount of liquid assets that the governing bodies “claim” to have. As the crypto market is not fully regulated, this claim cannot always be checked.

¹³ A very good analysis of the Terra LUNA crash can be found in (Liu et al., 2023). Interestingly, the authors argue that from the point of view of a fully rational player, as this is defined in terms of economic theory, this conclusion of the investor was not a sound one. That is far from saying that the ecosystem did not face inherent issues; it is just that the specific point of equal capitalization seemed to cause concern that were based mostly in psychological reasons and/or misunderstandings on the pricing of LUNA in relation to UST.

¹⁴ The DAO was called Anchor; it allowed UST deposits with high interests and many point this as one of the main reasons of the disaster. We will not go into details here since, even if that was true, our main goal is to investigate the runaway condition which was caused by the inherent stablecoin algorithm itself.

¹⁵ There were theories that the run was an orchestrated event. (Liu et al., 2023), provide solid arguments that this is not the case. However, even if there was an attack, this does not disprove that the Terra ecosystem had inherent security flaws, embedded in its core idea of algorithmic stablecoins.

Contrary to the case of *The DAO*, no mitigation was possible¹⁶. Many investors lost nearly all of their holdings that in some cases amounted to total of millions of dollars, with some of them even committing suicide as a result of the disaster¹⁷.

2.1.3 Protecting citizens in a DAO

It can be argued that the above two stories do not serve as an example of the dangers behind implementing and adopting a DAO, but are rather examples of the inherent greed and the get-rich-quick mentality that governs the actions of many cryptocurrency investors. However, this reasoning ignores that these disasters, apart from the intent of the key players, could only be possible due to the self-governed nature of the underlying DAOs. Fraud is indeed common in the crypto world, with many examples, such as the FTX or the Celsius Network case. However, in these cases, we had a traditional Ponzi scheme in which customer assets were used to draw new investors. The fact that these assets were cryptocurrencies is irrelevant. In the case of *The DAO* and the Terra LUNA disaster, however, things are different. Users were operating under the rules of a DAO, which proved inefficient.

The relevance of DAOs promoting open democracy is thus very evident. If the DAO fails to protect citizens' rights, how will citizens' trust be restored, and how will the damage be mitigated? Consider, for example, that a voting system is implemented, and a hack allows past votes to be disclosed. This damage cannot be undone. It is very important in this context to make sure that such breaches do not occur.

While a perfectly secure system is never achievable, we can build upon past experience to propose some mechanism that will further trust and minimize damage. Drawing again from the domain of economy, we can see that, despite the fact that the mechanism that initiated the crash is very specific to the crypto world and is difficult to happen in a regulated economy, bank runs generally are not that rare a phenomenon. Unlike the crypto DAOs, however, regulations can prevent bank runs from entering a runaway condition via a multitude of measures, such as government insurance of deposits, capital controls, etc. This paradigm can also apply to eGov DAOs by limiting the amount of self-governance that critical DAOs offer according to the use case at hand¹⁸. DAOs that delegate core governance functionality to experts, albeit in a transparent manner, can increase citizen trust and avert unwanted disasters.

On a similar route, the DAOs should ensure that no threatening transactions take place and that the governance model allows for no such transaction to take place. Threatening transactions in this context means transactions that compromise the privacy and safety of the data subjects.

¹⁶ Do Kwon, the creator of the Terra blockchain network created LUNA2, a second cryptocurrency and has offered some of its initial pool to previous holders of LUNA. The compensation that this action offered to the Terra LUNA disaster victims was not even marginal.

¹⁷ <https://www.investing.com/news/cryptocurrency-news/bitcoin-tenyear-prediction-michael-saylor-hints-at-price-boom-ahead-3493895>

¹⁸ It is obvious that if, at some point for example, DAOs are used for elections it would not be wise to allow citizens to vote upon the election mechanism. Technical aspects of the DAO moreover, such as what kind of algorithm is going to be used for encrypting votes, should not be part of the self-governance model.

DAOs used for elections, for example, should use custom encryption schemes such as homomorphic encryption to ensure that no information about voting preferences can be extrapolated from the blockchain. Conversely, DAOs that depend on transparency and an open ballot system should not allow the modification of rules to allow secret voting.

2.2 Markets in Crypto-Assets (MiCA) Regulation

The new MiCA Regulation is a new set of EU legislature that aims to provide transparency in cryptocurrencies in the EU and protect its subject against crypto-related fraud. At first glance, it is irrelevant to the aims of AI4Gov, and, indeed, MiCA is not explicitly considered in the framework or in any requirements that lead to the specification of application processes and components. However, as blockchain technology is adopted progressively in many e-government use cases and has the potential to be applied to public governance models in the future, MiCA could become relevant. The reason is that decentralized platforms that allow the interaction of political parties and candidates, lobbies, agencies, and citizens may progressively adopt mechanisms for allowing transactions that occur directly on-chain. These transactions may involve donations, renting services, etc. It goes without saying that such a setting if left unattended, is ripe for fraud. Malevolent parties may scam voters by falsely advertising donation wallets, vote-buying may become rampant, etc.

The businesses that MiCA regulates are termed Crypto-Asset Service Providers (CASPs), and these can be cryptocurrency exchanges and trading platforms, wallet custodians, and businesses offering crypto investment consultation services. Regarding the underlying assets, MiCA applies to stablecoins¹⁹ and utility tokens, which are cryptocurrencies that have limited usage in native infrastructure use cases, such as being used solely for transaction fees or used for in-app credit purchases²⁰.

MiCA dictates that all CASPs should be registered with a National Competent Authority (NCA) in a similar way that banks are obliged to do. The issuers of the coins covered by MiCA should always publish a white paper that provides information about the main details of the coin, such as information about the issuer, the coin issuing process, etc. While not all MiCA Titles have been finalized, MiCA already offers some protection mechanisms for citizens who use stable coins for transactions. In the AI4Gov scenarios, if we can imagine trading capabilities in future DAOs, the fact that MiCA limits the types of issuers and tokens means that a MiCA-compliant DAO will only involve assets that offer the required transparency, as this is dictated in MiCA. While shady transactions, like vote-buying, can again take place, the fact that the assets used are based on well-defined white papers and well-documented issuance processes makes tracking of such transactions easier to follow; this may limit the incentives of parties to engage in such business.

¹⁹ MiCA differentiates stablecoins to the so called “asset-referenced tokens” which are backed by commodities or a multitude of fiat currencies and to the so called “e-money tokens” which are coins pegged to a single fiat currency.

²⁰ The distinction between a utility token and a security token that can be used as money can often be difficult in practice. In many cases, there is nothing preventing trading of a utility token in the open market and making its price volatile and prone to speculation.

2.3 HyperLedger Aries

Allowing citizens to board the open decentralized platforms requires a mechanism for easy registration and verification of new users under the Verifiable Credentials Model that allows easy presentation and verification of evidence using a blockchain. The HyperLedger Fabric framework, which was documented in D3.1, is ideal for managing corporate and organizational blockchains; however, it can pose some limitations when used to accommodate wallets that belong to end users corresponding to physical persons. The main limitation is that each new user needs to be declared and registered, via their credential, to the blockchain network, which can be a burdensome process both for the citizen and for the admin organizations. Alternatively, a special chaincode can be deployed in peers that allows this registration to take place while also providing the necessary authorization and authentication mechanisms. This process involves the creation and maintenance of complex chaincode, which can be difficult to govern, especially when it involves cross-border cases, which may require custom authorization mechanisms depending on citizen nationality and/or location.

By adopting HyperLedger Aries, a framework specifically designed for identity management, many of these difficulties can be tackled. Aries supports the creation and administration of Decentralized Identifiers (DIDs) that can uniquely be resolved to each subject. Moreover, under Aries, the issuance and verification of credentials and attributes are straightforward. Typical use cases of Aries that are key to implementing mechanisms for citizen trust and participation include:

- Issuing of Self-Sovereign Identity (SSI) credentials which enable citizens to control their own identity.
- Evidence/credential verification mechanisms that allow wallet holders to present and provide proof of credential ownership (e.g., driving license) or attributes (e.g., legal age).
- Mechanisms for providing fine-tuned access to credentials based on the identity of the requesting party.

While the organizational smart contracts that govern the processes of policy-making are still implemented in the HLF infrastructure, the Aries framework is used for identity management at the citizen wallet level. The two solutions are linked naturally as Aries is a blockchain-agnostic framework; this allows the credentials that are verified at the Aries level to be used for performing smart contract operations at the HLF level.

2.4 Transparency of AI

For good or for worse, AI is now a key player in policymaking and a major actor in today's democracy. AI can be used for the creation of content that can then be used to manipulate decision making. Bias in AI, especially when the citizens are not aware of the underlying data sources that were used for training, can unfairly swing public opinion. Indeed, one of the AI4Gov project's main goals is to detect and point out such biases. With the advent of very complex LLMs, such as GPT, the AI can even impersonate behaviour and produce deep fake news that is extremely hard to detect, especially for the common person who is not trained in AI and prompt engineering.

Blockchain technology, due to its decentralized and transparent nature, can be used as a technology enabler to provide transparency in AI output. Since current LLMs are trained over a huge corpus of data and use millions or even billions of parameters for producing output, the 'straightforward' way of providing transparent AI by having each node validate the results of training and/or execution is, of course, not applicable. However, certain solutions, especially those inspired by the supply chain domain, can find applications.

One such example is the IPwe platform, which is an intellectual property transaction platform developed by IBM and is based on a synergy of AI and blockchain technologies. Similar to AI, the domain of IP involves complex relations and semantics that often involve cross-border legislation and cross-domain terminology that is often not fully aligned across legislation. AI leveraged NLP can help in identifying such patterns to generate summaries and reports for end users. Using blockchain, a shared registry of patents can be created, accompanied by the patent information extracted via AI.

Such approaches show the potential usages of AI in improving transparency rather than diluting it, provided, of course, that the AI can be trusted. It is exactly in this area of providing trust that AI4Gov can leverage blockchain to increase AI transparency. Two main techniques are relevant to this item:

- Anchoring the explainability report to any claim or document that is produced by AI. Although a claim may still be fake, producers who are confident in the truth of their statements can anchor explainability reports as further proof of their claims. This will help well-documented claims gain more visibility
- Using DIDs for the AI. By assigning to a generative AI an identity, the AI can behave as a physical actor in the blockchain. AI that has identity can be more easily trusted by citizens; this can create a trend in which AIs with identity are progressively preferred. The benefit of being able to discern the identity of an AI is that a consumer can check its parameters such as training algorithm, past predictions, previous bias reports etc.

3 Decentralisation in AI4Gov

This section corresponds to Section 3 of the first iteration, D3.1; it should be viewed as the second and final version of this Section. The bulk of data requirements remains the same, with the exception of data required for citizen participation. Thus, the data requirement tables will be repeated here for completeness, together with the extra requirements needed.

3.1 Data information in AI4Gov

This updated view for data requirements is depicted in Table 1, where data and end-user information are grouped per pilot.

Table 1: Types of data and end users per pilot case

Use Case	Pilot	Data	Type	Users
Water Management drinking water –	DPB	-Sewage Treatment data	-Static/Streaming -Policy reports	-Workers at the municipal consortium for water management -Local administration -Citizens
Water Management sewage water –		-Water cycling billing data -Streaming sensor data -Policy data -Citizen wallet data -Explainability reports -Bias reports		
IRCAI global 100 projects	JSI	-IRCAI data of projects submitted (textual description, URLs) -Event Registry data (news and event items)	Static	-Teams in private or public Institutions/Organizations that are submitting projects to the IRCAI Global Top 100 program. -Government -Corporate -Researchers
SDG Observatory				
OECD policy document analysis				

		-OECD AI policy initiatives -Explainability reports -Bias reports		
Parking tickets monitoring	VVV	-Census data -Household water data	-Static -Policy reports	-Policy makers -Citizens
Waste management – Pay as you Throw		-Tourist data (arrivals, overnight stay, cruise data) -Airport traffic data -Municipality events attendance data -Citizen wallet data -Explainability reports -Bias reports		

Most data requirements remain the same, with the difference that now that citizens are involved, the relevant data should also be included. Most specifically, we extend the use cases of the DPB and VVV pilots to allow for the possibility of citizens being part of the blockchain. These citizens will have access to policy data that represent the policies that these two pilots define and implement. Two DAOs will be implemented as part of the Policy Recommendation Toolkit; one will be used to propose and vote policies for DPB and one for VVV. While the JSI pilot remains the same, from a data requirements point of view, it will support the two DAOs with a suite of smart contracts; these smart contracts will allow the anchoring of OECD and IRCAI reports to the policies that have some relevance. The anchoring will be done by peers that belong to DPB and VVV, and they will be available for review by citizens who use the citizen wallets that these two organizations distribute.

The citizen wallet that is implemented naturally holds information about the citizen, such as credentials and proofs requested by the citizen. These are not stored centrally but are fully

controlled by the citizen; they are listed here for completeness to make clear that such data are involved in the extended use cases that allow citizen participation.

3.2 Decentralised Data Storage in the AI4Gov platform

Figure 3 depicts the overall AI4Gov reference architecture as described in D2.3. Adding to what was documented in D3.1, the blockchain architecture will now accommodate citizen wallets and DAOs. Citizen wallets are not very different from the original organizational AI4Gov in terms of how they interact with the other components; however, the implementation details are different since they depend on identity management based on HyperLedger Aries. This will be explicitly shown in the various layers of the architecture.

DAOs on the other hand consist of smart contracts, which already were embedded in the architecture in D3.1. However, the more decentralized governance model has to be reflected both at the business and at the application level. Instead of defining blockchain governance model via configuration of the channels, we now have an open interface which can be used for blockchain governance operation by all participants.

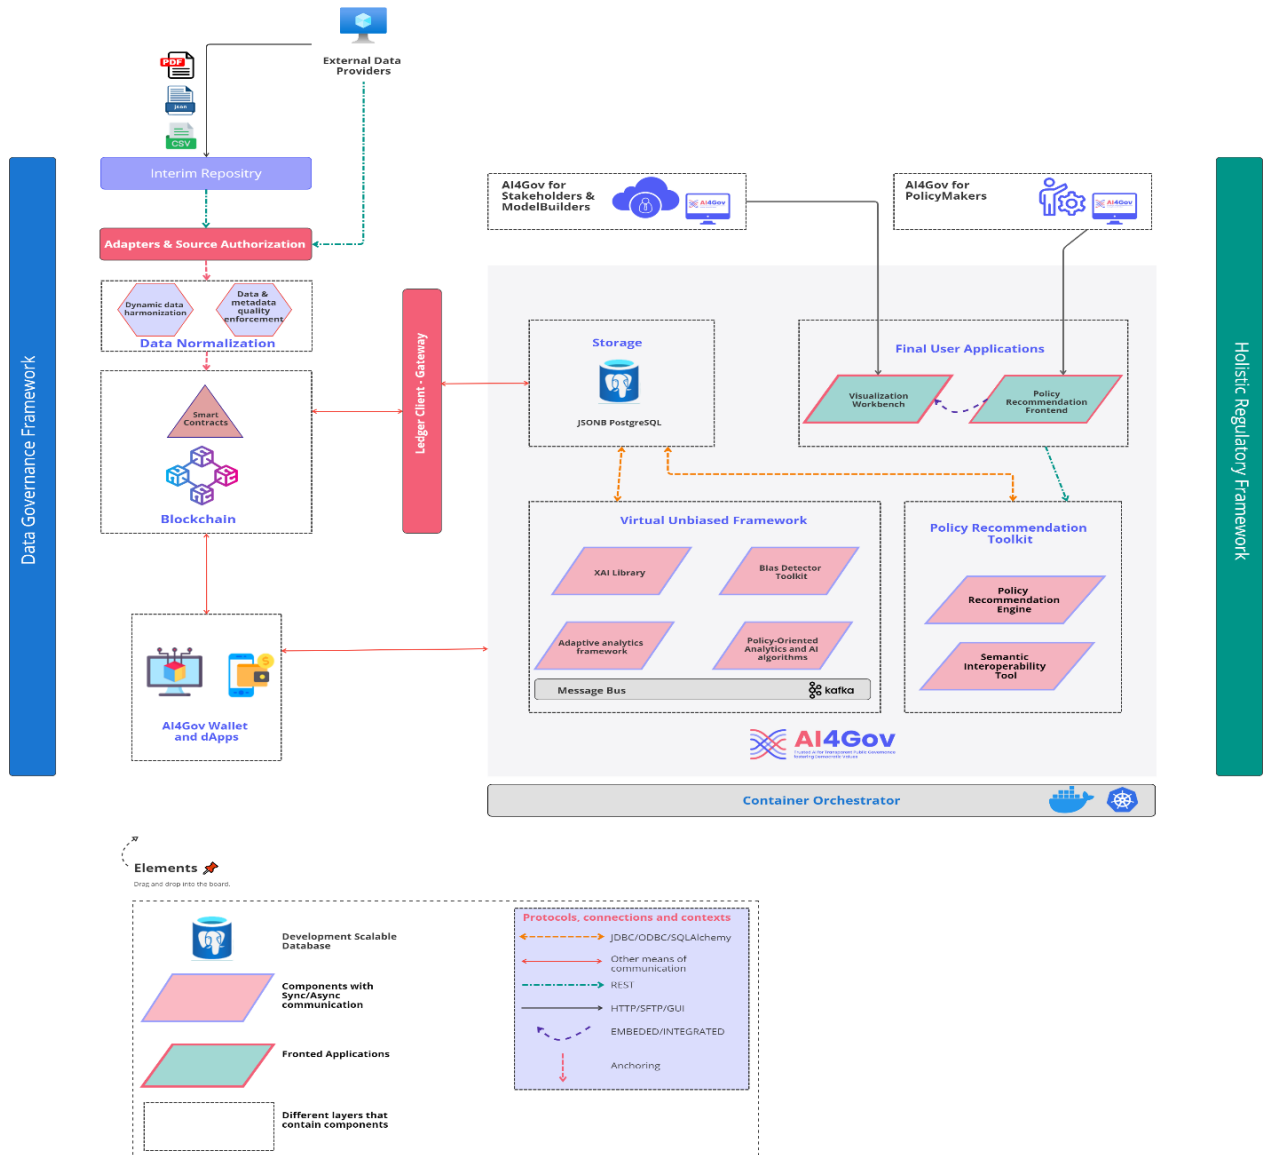


Figure 3: AI4Gov Reference Architecture

One of the main points of the architecture left open in D3.1 reads: “While not clear at the present moment, a decentralized data storage mechanism and data validation mechanism may also be beneficial for end users by, for example, anchoring the explainable data to record the rationale for a result produced at a specific point in time.” We can now say with confidence that such storage and validation mechanisms are beneficial and in fact central for facilitating open democracy.

Concerning open democracy in particular and its relation to requirements and architecture, we must emphasize that the modified data requirements listed in Table 1 serve only as a staging point for defining the mechanisms for open democracy. It is a straightforward way to expand the

definition of pilots to include citizens in policymaking. However, the vision of the architecture is broader; it is to support citizen engagement in any open process by giving the means to implement and run the relevant DAOs for each use case. The mechanisms implemented for citizen participation for the extended pilots should be easily adapted to any case that involves citizens. The redesigned architecture is targeted exactly towards this goal.

3.3 Architecture for Decentralized Data Governance

The second and final iteration of the architecture involves the finalization of the technical layer that describes the application and infrastructure components. As the architecture was refactored to accommodate for the new functionality, we will list the previous diagrams to highlight the modifications. As was the case in D3.1, each layer will be presented separately.

3.3.1 Business layer

The first iteration of the business layer can be seen in Figure 4; the reader can refer to Section 3.3.1 of D3.1 for the explanation of the various components. Since the solution has now expanded to include a whole new category of stakeholders, namely that of citizens, the motivational viewpoint is now presented separately in Figure 5. Here, the “blockchain value” element of the original architecture has been renamed to the more appropriate “Additional Value” as a broad term to include all new value that is offered. The value elements remain the same with the addition of the “inclusiveness” value. These values are associated with both groups of stakeholders, citizens, and governance agencies. One of the goals of governing bodies is to conduct Efficient Policymaking. This goal is influenced by the added values and is also realized by the Transparency principle. The other two goals are closely related. They are “Engage citizens” for government stakeholders and “Participate in Governance” for citizen stakeholders. These are realized by two principles: Transparency and Open Democracy. Two requirements for the participation of citizens are the ability to “Assess programs”, meaning the various political platforms that are presented publicly and to “Assess news items” in a trustworthy manner.

With these in mind, the business viewpoint is amended, as depicted in Figure 6. Here, the original “User” group that was designated for an organizational user having any role of those identified in D6.1 has been separated by the Citizen user. The business interface that depicted the special process of accessing a decentralized platform has been removed and is now understood to be “merged” with the two service components that serve the two user groups. The original “User” is served by the Decentralized data governance service, and the Citizen is served by the Open Democracy via decentralization service. These two services provide an extra layer of abstraction between the stakeholders and the business product. Whereas in the first iteration, users were directly associated with the “AI4Gov Decentralized Infrastructure and Contracts” product, we now have an extra layer of business processes that describe what processes are followed to realize the respective services.

More specifically, the original “provide decentralization” service that was explicit for organizations, is now renamed “Decentralized Data Governance.” This service is, in turn, realized by the “Efficient Policy Making” business process. The “Open Democracy via decentralization”

service is likewise realized by the “Assessment of platforms” and the “Decentralized voting” business processes; the first one describes the evaluation of platforms, policies, and programs, while the second one describes the process of actual voting. Both are aggregated by business functions, such as “propose draft,” “access bias reports,” etc.

The product is renamed “AI4Gov Decentralized Infrastructure and DAOs” to indicate that the platform now supports fully functional DAOs instead of smart contract suites. Since the services have been moved to the upper layer, the product now consists of two business processes that implement the product. It is to be noted that this relation is a “realize” relation and not composition or aggregation. The diagram should read: “DAO governance realizes AI4Gov Decentralized Infrastructure and DAOs” and “Decentralized data processing realizes AI4Gov Decentralized Infrastructure and DAOs”. The original contract relations remain the same; the only difference is that the “Decentralized business logic agreement” has been renamed to the more general “Decentralized DAO logic agreement.”

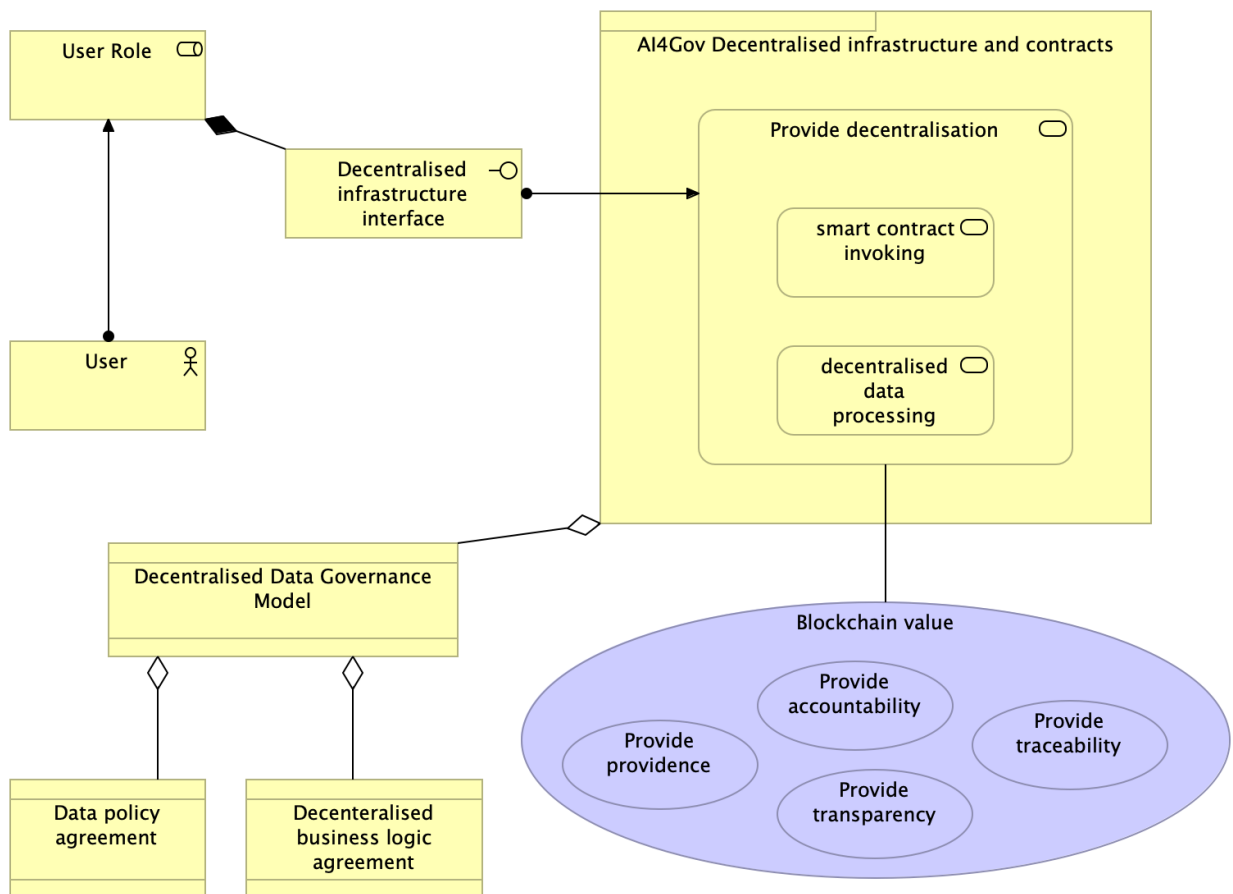


Figure 4: 1st iteration of the Business layer of the decentralised architecture

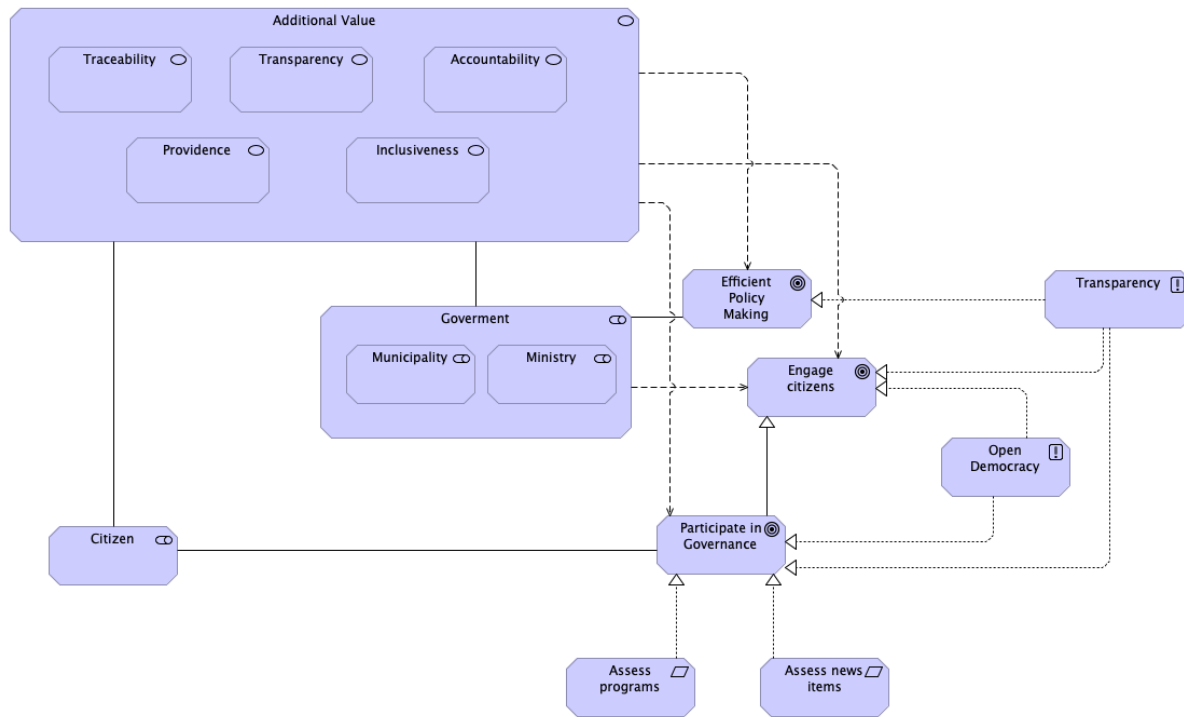


Figure 5: Motivational viewpoint for the decentralized architecture

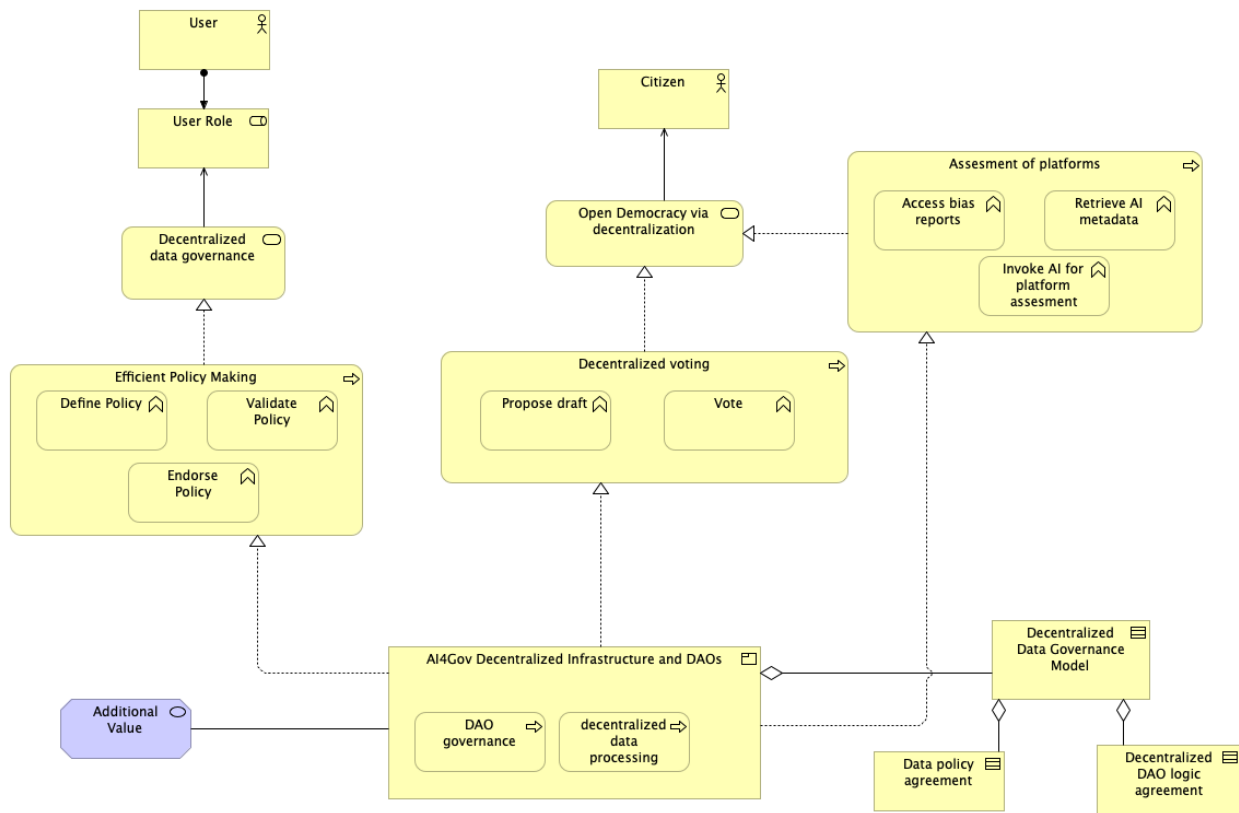


Figure 6: Business viewpoint of the decentralized infrastructure

3.3.2 Application layer

The original Application viewpoint of the 1st iteration is depicted in Figure 7. In this iteration, the application layer has been amended in the following ways:

- It was refactored to realize the needs of the new business viewpoint that involves citizens and DAOs.
- It involves the instantiation of Architectural Building Blocks (ABBs) to Solution Building Blocks (SBBs) to accommodate for the fact that the technical requirements are now fully derived and a full prototype has already been deployed.

As such, the second iteration of the Application layer is very different from the first one, which is depicted here mainly for reference. Therefore, there will be little effort to explain how the diagram changed, and the present section will mainly focus on explaining the viewpoints. There are two viewpoints. The first one is the ABB viewpoint, which shows what the application functions and processes are expected to do to support the business viewpoint. The second one is the SBB viewpoint, which describes *how* these functionalities and processes were implemented in AI4Gov.

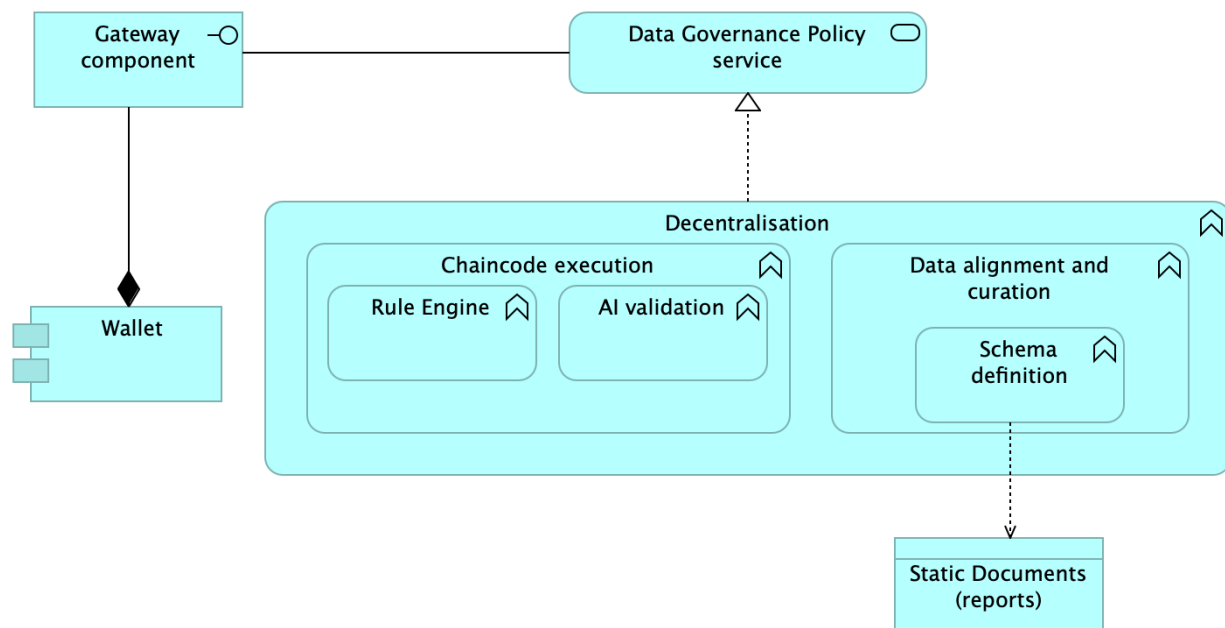


Figure 7: 1st version of the Application layer of the decentralised architecture

3.3.3 ABB viewpoint

The ABB viewpoint is depicted in Figure 8; in order to avoid the already big clutter, several internal relations have been omitted. The ABB viewpoint shows how the application functions and processes realize or serve the main business processes that fuel the services offered to the AI4Gov stakeholders. One core application function is the “Decentralization” function, similar to the

original “Decentralization” function described in D3.1. This is composed of two application functionalities: the DAO functionality, which is now expanded to include all smart contract definition and invocation functionalities, and the original data alignment functionality for aligning and anchoring data on the blockchain. Naturally, the “Decentralization” function serves all other application processes. The Identity Management and VCs application process consists of all the modules that involve the identity of users, as well as the core mechanism for presenting and verifying verifiable credentials. The Identity Management process is served by the decentralization function, and it serves the remaining application processes.

The remaining application processes, as can be seen in the Figure, realize or serve the various business aspects of the solution. The Policy governance and Policy administration application processes for example, consist of application functions that realize the various business functions of the “Efficient Policy Making” business process. The realization/serving relations may be one to one or one to many; the important point is that each business function is realized by a corresponding set of application functions or processes.

The more complex part of the ABB viewpoint involves the assessment of content. This involves many application functions that may operate alone or in synergy to realize one or more business functions that are part of the “Assessment of platforms” business process. The anchoring of reports to content is a key enabler for this process as it allows end users to retrieve and check the off-chain reports that accompany the content. Other functions include the validation of the reports and/or the metadata, which checks if the hash is still valid, meaning that the report has not been removed or modified off-chain. In case a report is accompanied by a model, this can also be run on-chain or off-chain to reproduce the results.

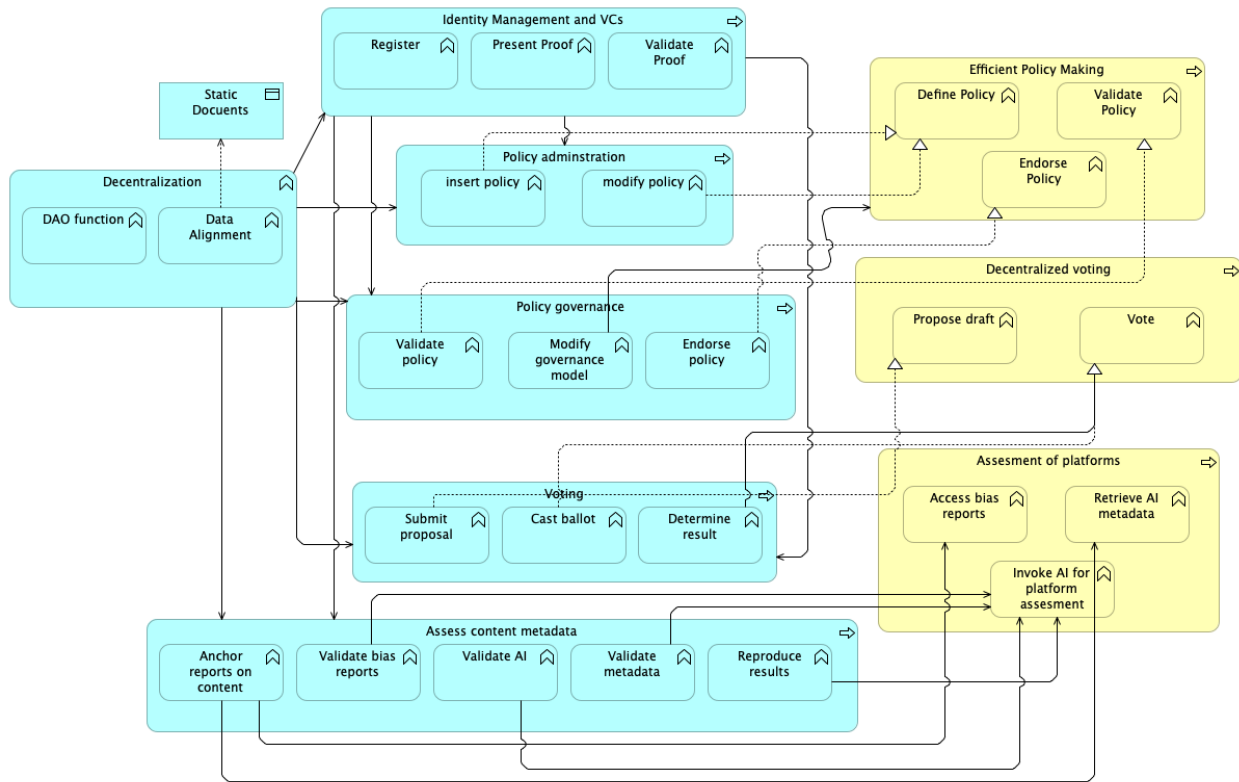


Figure 8: ABB viewpoint of the decentralized architecture

3.3.4 SBB viewpoint

The SBB viewpoints consist of all the software and hardware components that materialize the application functionality that is described in the ABB viewport. There are various approaches in designing an SBB viewport. We followed the approach of assigning components to application functions and processes; when a straight “assignment” relation is used, that means that the component basically implements all functionality of the process/function. More granular relations are depicted between subcomponents and processes/functions.

Unfortunately, the viewport is too big to be displayed in a single diagram. Therefore it will be separated into a set of smaller ones, each one showing how the relevant components implement various functionality.

We start with the core functionality that includes the “decentralization” functionality itself and the identity management. The relevant SBB diagram is depicted in Figure 9.

The BIE component refers to the whole blockchain infrastructure. This infrastructure hosts two frameworks, the HyperLedger Fabric, and the HyperLedger Aries. As a whole, the BIE component is assigned to the Decentralization function. AI4Gov hosts a number of smart contracts on the HLF infrastructure in the form of a chaincode. One chaincode implements the PolicyDAO component, which is the DAO implemented in AI4Gov, to demonstrate the open democracy model using the blockchain, under the scenario of citizen co-creation in policy making. As such, the DAO contains

all mechanisms for proposing policies, providing feedback, and voting. It realizes the DAO subfunction of the Decentralization function; note should be taken here, however, that, in general, there could be multiple DAOs realizing different open democracy scenarios in the general case. The Data Anchoring smart contract is the set of chaincode that has been implemented to allow the insertion of new blocks and anchoring of off-chain files to the blockchain. The HyperLedger Aries infrastructure, lastly, is the component that works behind the scenes to register new users, validate their presented proof, and, conversely, present them with certified proof.

End users access this functionality via their wallets, an organizational one for organizations and a citizen one for citizens. The architecture supports the creation of a wallet by an organization via the Aries framework as well if this is better suited to the organization's needs²¹.

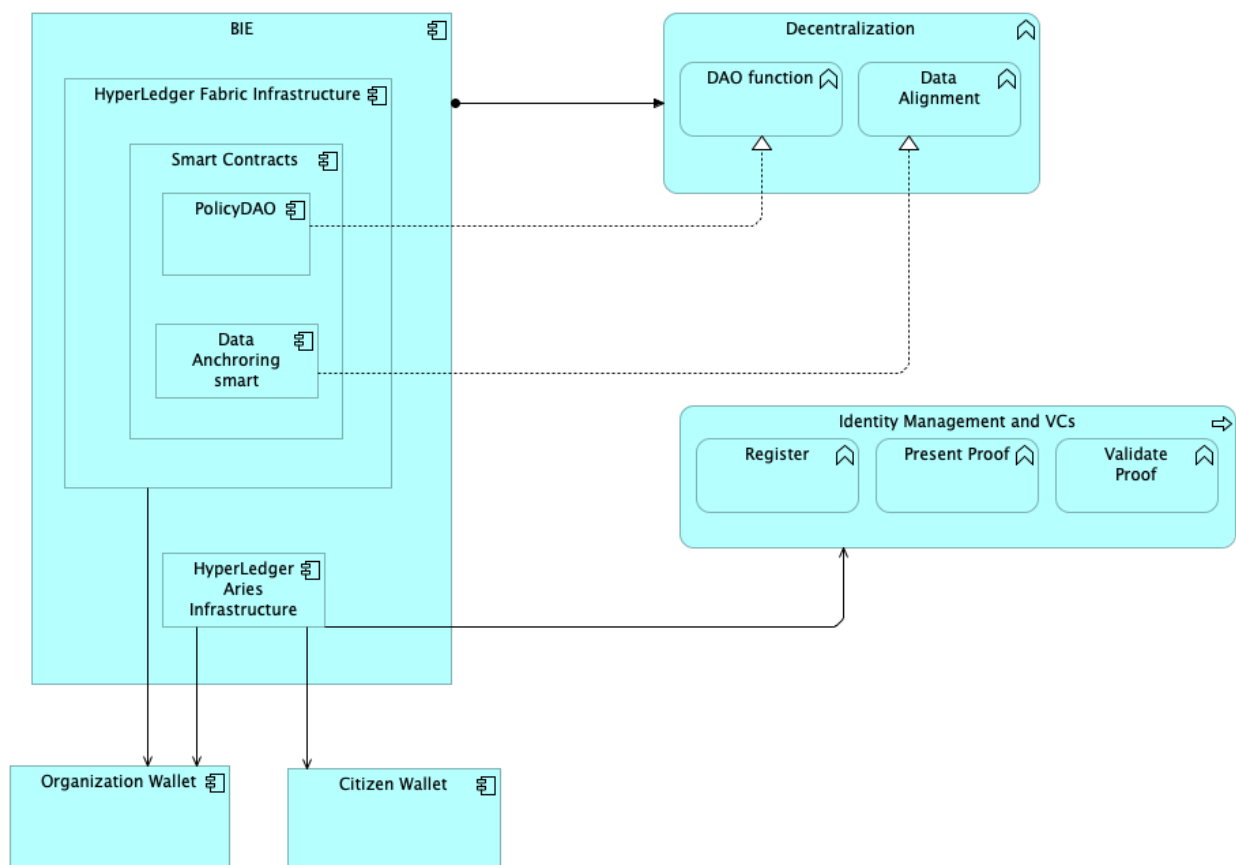


Figure 9: SBB diagram for the Decentralization function and the Identity Management and VCs process.

²¹ For example, if an organization boards the blockchain mainly to participate in open democracy processes and is not interested in executing organization exclusive smart contracts, it can board the platform via a citizens' wallet. Since it will present their certificates upon registration, their role should be clearly assessed and therefore properly handled by any smart contract.

Figure 10 depicts the SBB diagram for all processes that have to do with policymaking, both from the organizational and from the citizen aspect. As was the case depicted in D3.1 and D3.3, organizations can directly insert and modify policies in the blockchain in a transparent manner, by which the evolution of a policy maintained by an organization can be viewed in the blockchain. This is the “policy administration” process that is realized by the “Data Anchoring smart contracts,” which anchor any off-chain data that accompanies the policy, and by the PRT smart contracts, which are the smart contracts that actually insert and modify the policy data on-chain.

The “policy governance” process is likewise realized by the PRT chaincode. Although not depicted in the diagram for clarity, this component has a number of sub-components for endorsing a new policy, thereby advertising it to other organizations, for validating a policy by checking via smart contract code if it accomplishes certain KPIs and for modifying the governance model which defines how a new policy is inserted and/or modified²².

The part that bridges organization governance with citizens is achieved by the “voting” process, which is mainly realized by the PolicyDAO component. Again, not all sub-components are listed for clarity, but the PolicyDAO allows citizens to invoke smart contracts that mine a draft proposal in the blockchain, voting mechanisms for submitted drafts, and algorithms to calculate the result of the vote²³. All these can, of course, be audited by any party at any time using the core functionality that the BIE component offers.

²² In the current implementation a single organization can insert a policy that it proposes or adopts, but the governance model allows this to change and only allow new policies to be inserted by, for example only specific organizations or after a certain number of organizations have endorsed the new policy. The way this mechanism was implemented and can be configured via HLF, is explained in D3.1

²³ Preserving the secrecy of the vote involves the invocation of special Zero Knowledge Proof (ZKP) protocols. These are implemented in the context of the Policy Recommendation Toolkit and are explained in D3.3. Currently, they are not yet incorporated into the PolicyDAO, which means that the DAO does not preserve vote secrecy. By the time that D3.4 is composed, it is expected that the ZKP functionality will be migrated into the blockchain’s DAO by appropriate smart contracts.

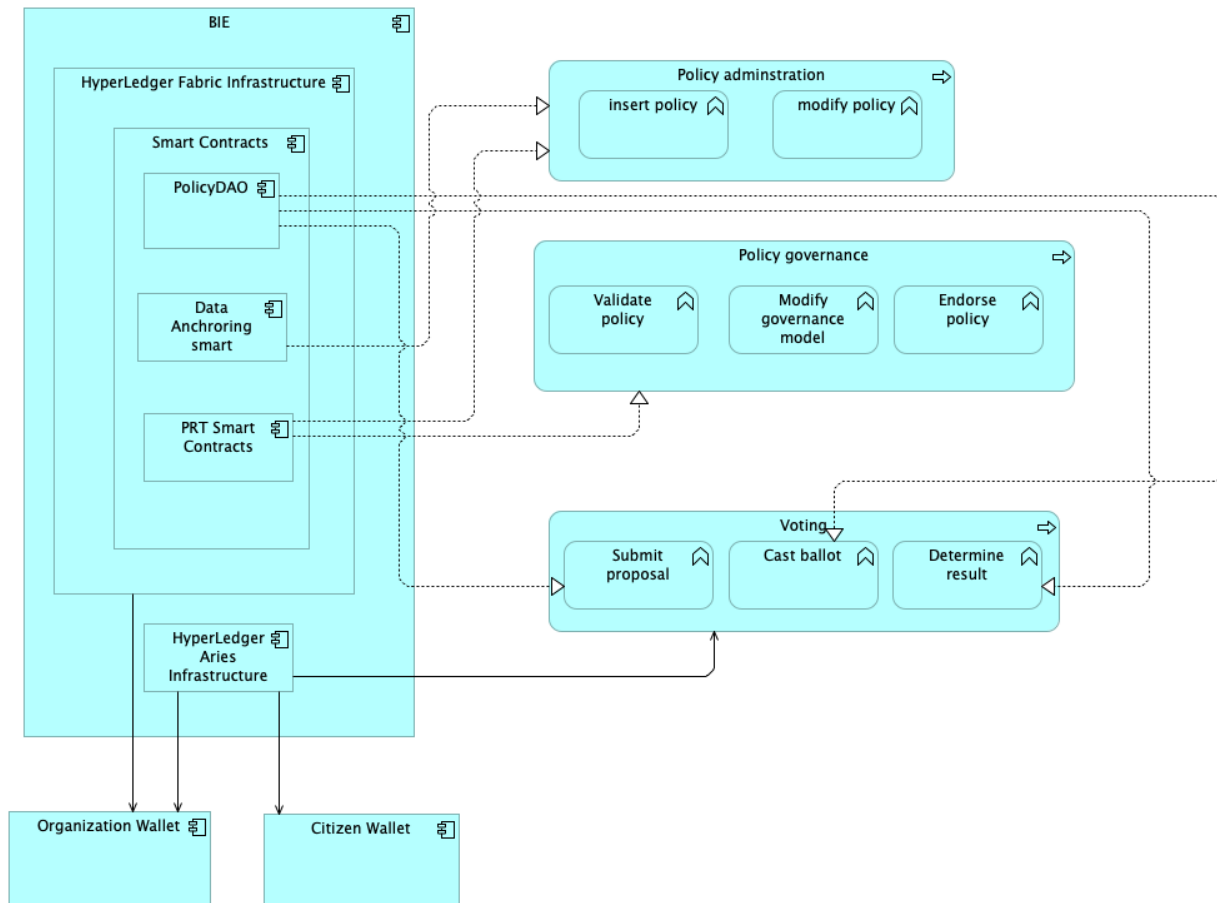


Figure 10: SBB diagram policy administration functionality

Finally, Figure 11 depicts the SBB diagram for the component achieving the assessment of content. The “Data anchoring smart contracts” insert report summaries or anchor large reports on the blockchain using the mechanism described in D3.1. These anchors produce hashes that are stored on-chain. The Data Validation smart contract checks for the validity of this hash; if the hash of the off-chain file is not the same as the one stored in the blockchain, that means that the report has been tampered with. The AI validation smart contract refers to a set of components that can use AI to check for spurious input or validate results. Strictly speaking, it is not an SBB since, by the time of the writing of this document, it has not been fully implemented. Currently, there is only an experimental instantiation of the component that checks if the report of a simple AI model works as advertised. This instantiation works as follows:

- The user provides an instance of an AI model via a file, namely the algorithm of the model and the values of weights. Since the model is to be validated by all nodes of the blockchain, its output should be deterministic. The models currently supported can be instances of the Support Vector Machine algorithm and the Polynomial regression classification algorithm.
- The user provides an input and an output and asks the component if the algorithm predicts the same output.

Although of limited practical use, this component shows the feasibility of anchoring models and executing them on the blockchain. Using the functionality that this component offers, more complex scenarios will be constructed and demonstrated in D3.4.

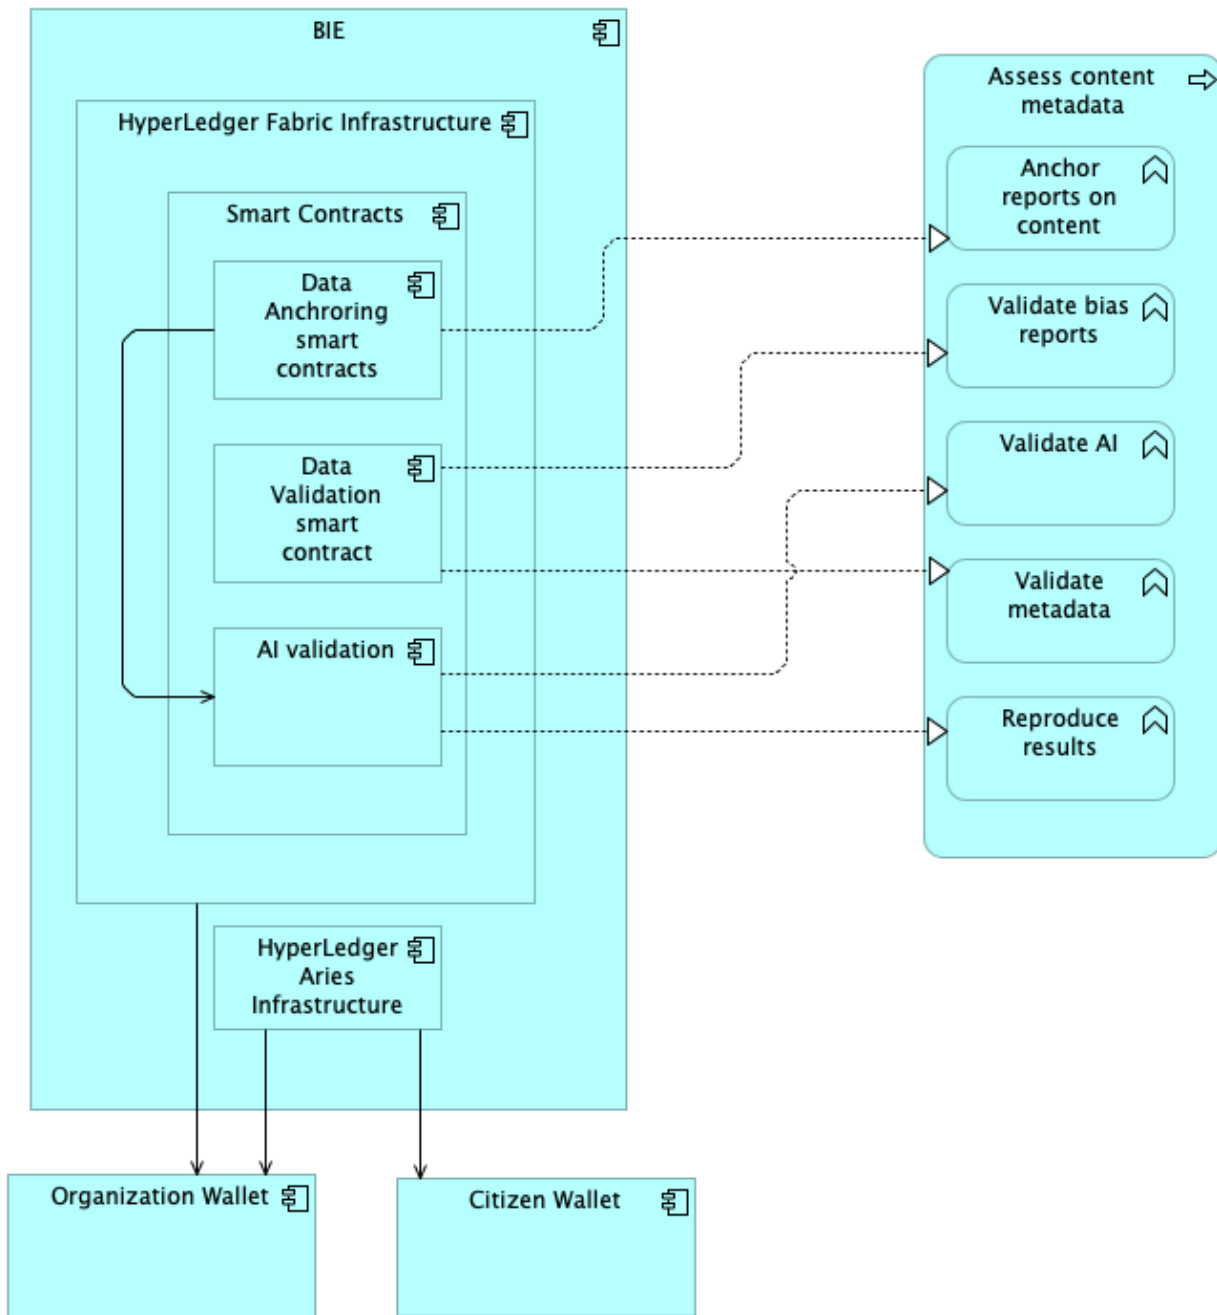


Figure 11: SBB diagram for content assessment

Lastly, the technical infrastructure aspect of the decentralized data framework is depicted in Figure 12. The deployment of resources such as VMs and networks for the full AI4Gov is described in D2.3. Here, we depict only the relation between the physical components of the BIE and how

these serve the different components. Briefly, a set of VMs was used to instantiate a test blockchain network. The blockchain has three peers corresponding to the three pilots plus a node that runs the orderer service that achieves consensus in the HyperLedger Fabric setting. Two additional nodes have been instantiated for the purposes of demonstrating how the Policy Recommendation Toolkit can be used together with the BIE to achieve transparent Policy Making. The "Generic Municipality" node represents a municipality that can, apart from creating and modifying policies, endorse existing ones so that they are advertised to governance agencies. The "Generic Ministry" node represents a ministry that can view popular policies and then tag them for consideration in lawmaking.

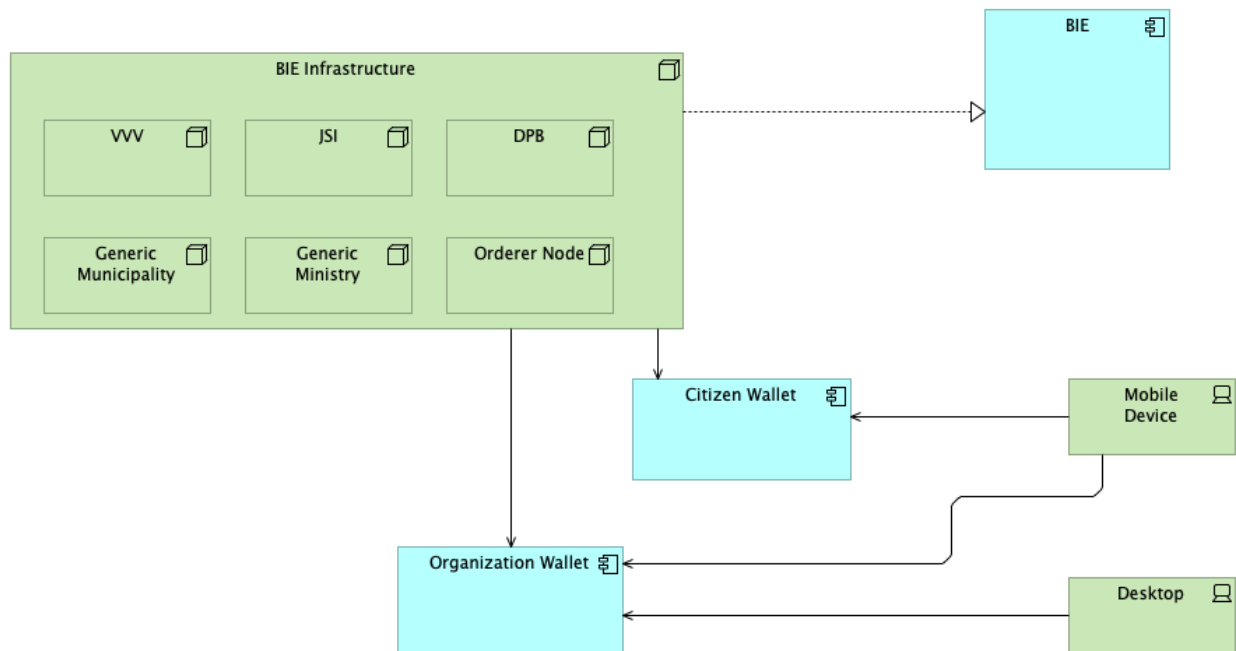


Figure 12:SBB for the technical infrastructure viewpoint

3.3.5 General requirements

The general requirements regarding the data governance for each type of data and actor were defined in D3.1 and are reiterated for completeness in Table 2, together with any changes proposed. At first glance, we should expect that the new requirement for citizen participation should produce major amendments in the data governance model, but this is not the case. The original table was formulated in terms of data owners, and as such, the same conditions apply, provided that we also include citizens in the data owners group. For example, a citizen can anchor, version, and delete data on the blockchain just like an organization. Similarly, the citizen can maintain chaincode if it is part of a DAO that is part of. The only difference here is that in the case of DAOs, we might want to differentiate voting rights by role. In a policy voting DAO, for example, changes in the voting process may be subject to change only by voting between constitutional

peers. Table 2 was amended to denote that voting that has to do with smart contract maintenance is now role-based.

Policy description	Type	Location	Policy(ies) adopted
Maintenance and upgrades of the blockchain source code	Source code	Off-chain	External governance
Read access of anchored data	Data	Off-chain and On-chain	Defined by data's owner
Insertion of anchored data	Data	Off-chain and On-chain	Data owner
Versioning of anchored data	Data	Off-chain and On-chain	Data owner
Deletion of anchored data	Data	Off-chain	Data owner
Maintenance of smart contracts code	Smart contract code	On-chain	<ul style="list-style-type: none"> • Single node • Majority role-based vote • Minimum number of role-based endorsements • Unanimous role-based vote
Validation of smart contract invocation results	Smart contract code	On-chain	<ul style="list-style-type: none"> • Single node • Majority role-based vote • Minimum number of role-based endorsements • Unanimous role-based vote

Table 2: Decentralised Data Governance policies in AI4Gov

4 Technological enablers

In D3.1 we gave a brief overview of the main technology enablers that facilitated the implementation of the decentralized data governance framework of AI4Gov. These were the HyperLedger Fabric framework, which provides a private blockchain solution that allows the creation and evocation of smart contracts together with mechanisms for governing data and code via a custom set of on-chain and off-chain rules, and the OpenDSU framework for the implementation of end-user wallets.

The reasons for adopting HLF still stands in the final iteration of the document. The technology will not be reiterated here; the reader can refer to Section 4.1 and the relevant subsections of D3.1 to review the core characteristics of the technology.

The potential use of HyperLedger Aries to implement ZKP mechanisms was also mentioned in D3.1; under the new requirements of designing solutions involving citizens and citizen co-creation, it was decided to fully adopt HyperLedger Aries for the citizen wallet.

Regarding OpenDSU, the fact that it showed a slow performance under experimental settings raised considerations of how it would scale in a production environment. Although a blockchain such as the EBSI offers much more powerful nodes than those used for the AI4Gov test environment and can thus, in theory, make an OpenDSU-based wallet run smoothly, since a new wallet of the citizens was, in any case, decided to be implemented from scratch, it was decided to abandon the OpenDSU technology and redesign the organizational wallets.

In the next subsection we are going to provide a brief overview of the HyperLedger Aries framework and how this can enable the identity management needed for citizen participation in AI4Gov.

4.1 The HyperLedger Aries

HyperLedger Aries is a blockchain-agnostic library suite that uses ZKP primitive instructions to offer verifiable credential functionality. It can be used to create blockchain interface layers, also called resolvers, for initiating and signing blockchain transactions. With Aries, communication can be performed both on-chain and off-chain via appropriate messaging systems. Aries includes wallets that allow secure storage of keypairs as well as APIs for supporting higher level protocols.

A high-level diagram depicting the main functionalities of Aries is depicted in Figure 13. Aries builds upon the Ursa Crypto library, which is a library developed by the HyperLedger Foundation that implements cryptographic primitives²⁴. Central to Aries lies the concept of the so-called HyperLedger Aries agent. An agent is a component that is designed to manage the exchange of digital identity in a secure manner. The agent operates on the client side, and it interfaces with the blockchain only when this is needed to resolve the DID. By default, the HyperLedger Indy is used as a resolver, but other resolvers can be used as well.

²⁴ Ursa is now in the end-of-life status and its components are migrated to the relevant HyperLedger projects.

Hyperledger as a Verifiable Information Exchange Platform

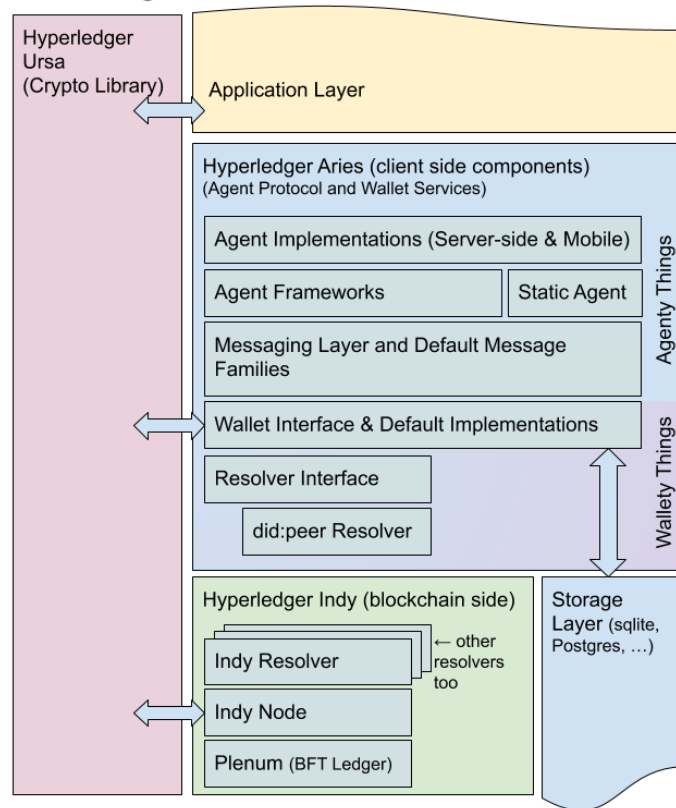


Figure 13: The HyperLedger Aries stack

4.2 Prototypical implementation

The current section will provide a brief presentation of the status of the implementation of the decentralized data governance framework. As this is the final iteration of the “Decentralized Data Governance, Provenance and Reliability” deliverable, the prototype provides nearly the full capability of the decentralized data governance framework²⁵. It is to be noted that, as a standalone component, the framework can have a more generic functionality than that realized by the instantiations of its ABBs. In the context of the project, these instantiations involve a framework for handling data and smart contracts involving policy recommendations and a framework that allows citizen participation in the policy recommendation process. However, the same framework can be used for other use cases involving organizations and/or citizens. For

²⁵ There is some experimental work being performed for the AI validation component as was mentioned in Section 3.3.4, closely related to T3.3.

example, the certification of citizens can involve university credentials and the voting DAO be substituted by a DAO enabling participation in open campus processes.

4.2.1 Decentralized Policy Making

This scenario demonstrates the main functionalities that have to do with data governance from the organizations, together with the smart contract capabilities that allow custom policy making and policy recommendation based on business rules encoded in smart contracts. The scenario starts with assuming the identity of a user belonging to VVV who uses the Policy Recommendation Toolkit, which is integrated into the Visualization Workbench (Figure 14). It should be noted that although the user is logged in via the Visualization Workbench, the VVV user is also defined in the HyperLedger Fabric test network, and the user is fully verified in the blockchain as well. This is important as only identified HLF users can execute smart contracts.

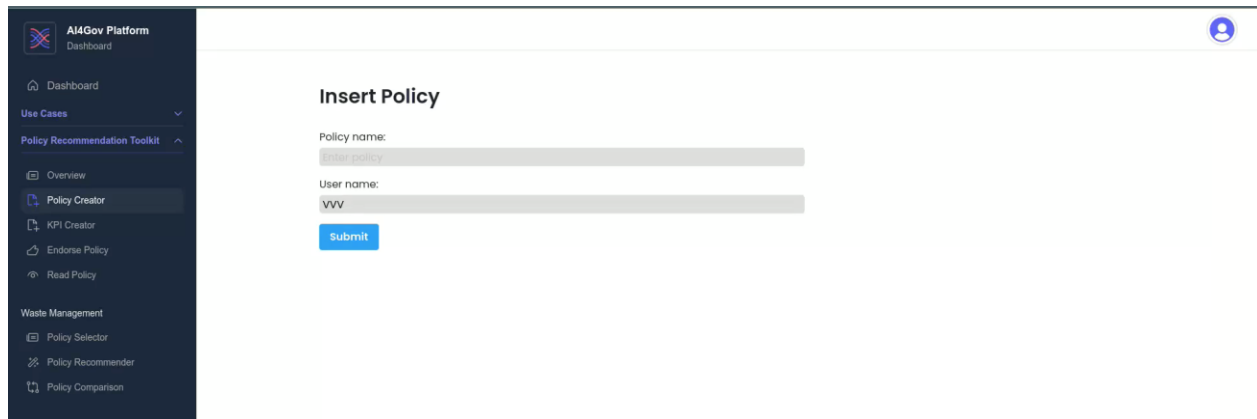


Figure 14: PRT – Insert Policy Screen.

The initial screen for policy creation allows the user to insert a new blank policy by giving its name. The user gives the name, and after some time, the policy is inserted into the blockchain. As with everything in HLF, this transaction that changes the status of the chain is done via smart contracts. The current governance model of the HLF that was implemented for the needs of the scenario dictates that any user can insert policies in the blockchain. This can change and demand a custom rule that demands a certain number of blockchain endorsements to occur before an entry is made. This fact highlights that even seemingly “simple” operations, such as data insertions, can be governed by custom rules.

AI4Gov Platform
Dashboard

Dashboard

Use Cases

Policy Recommendation Toolkit

Overview

Policy Creator

KPI Creator

Endorse Policy

Read Policy

Waste Management

Policy Selector

Policy Recommender

Policy Comparison

Insert Policy

Policy name:
Waste management Policy

User name:
VVV

Submit

Success! Policy inserted

Figure 15: Policy created successfully.

Continuing with the example, the user can define a set of KPIs that the policy achieves. These are two inserted into the blockchain. Behind the scenes a dictionary was created to act as index of the inserted KPIs so that they can be referenced by other policies. The index was anchored off-chain and is updated with each new KPI to act as a taxonomy for KPIs.

AI4Gov Platform
Dashboard

Dashboard

Use Cases

Policy Recommendation Toolkit

Overview

Policy Creator

KPI Creator

Endorse Policy

Read Policy

Waste Management

Policy Selector

Policy Recommender

Policy Comparison

Insert KPI

Enter KPI name & value:

Select Policy:
Waste management Policy

Selected option: Waste management Policy

KPI name	KPI value
Reduction of fuel costs of garbage truck	20
Reduction of time to collect waste from	20
Increase of green and organic waste pr	40

Submit

Success! KPI inserted

Figure 16: Association of KPIs to a policy

In the next step, we assume that the user wishes to retrieve a list of recommended policies based on the fulfillment of some target KPIs. The user selects a set of hard constraints that should be mandatory achieved and a set of soft constraints by which the returned policies will be ranked. All constraints are accompanied by a target value and the condition that the constraint should fulfill (i.e., greater, lesser, or equal to), as is depicted in Figure 17.

Figure 17: Defining criteria for policy recommendations

The tool returns the policies that fulfil the criteria together with the ranking. It does so by invoking a smart contract that filters the policies and computes their score. As is evident by the output, the smart contract ranks the policies based on the percentage of the soft KPIs that they fulfil, with a score of 100, meaning that the policy fulfils all soft constraints. Hard constraints do not enter into the score evaluation but are only used for policy filtering. This is just the way the smart contract was implemented for the purposes of demonstration. The infrastructure allows the modification of the smart contract logic to provide more custom or different output (e.g., by also scoring the hard constraints). This is supported by the decentralized data framework by employing directly the relevant mechanisms of the HyperLedger Fabric chaincode lifecycle management. A peer may change the code and send the change for approval to the blockchain. If enough peers vote for the change, according to the rules governing the chaincode governance, the update is approved and can be committed to the blockchain by any peer.

Figure 18: PRT recommendations using a smart contract.

4.2.2 Identity Management and Verifiable Credentials

For the second scenario, we start with an empty citizen wallet that holds no credentials (Figure 19).

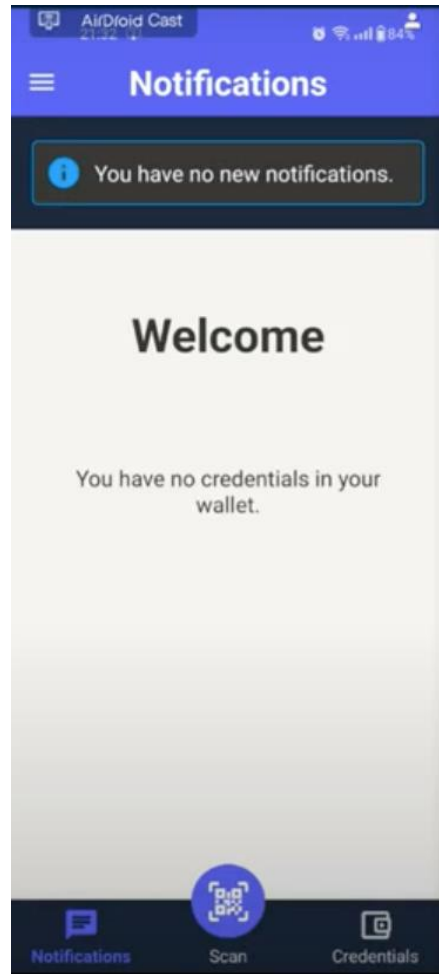


Figure 19: Citizens' wallet – Initial screen

An organization that distributes the wallet, such as the government or a municipality, can generate invitations to agents, thus allowing them to board the platform. One such invitation captured by the backend of the Aries infrastructure is depicted in Figure 20. This QR code is depicted in the backend for demonstration purposes. In a real-world scenario, a citizen would make a request to join and the organization would issue the invitation and present it in a relevant environment, like in a push notification or in another form of message.

The user can now accept the invitation by scanning the transmitted QR code (Figure 20 left); after some time, she/he is connected with the issuer of the invitation, which, in this scenario, is VVV (Figure 21 right). After the invitation, the issuer can issue a full credential and offer it to the citizen (Figure 22 left); if the citizen accepts, she/he now has a credential filled with all the attributes sent by the issuer (Figure 22 right). The citizen can verify that the credential presented in her/his screen has the same is the same as the one recorded in the blockchain; she/he is free to reject

the credential, if a mismatch is identified. This credential can now be presented to any party requiring proof under the VC scheme²⁶.



Figure 20: Boarding invitation generated by Hyperledger Aries

²⁶ A common misunderstanding is that this scheme proves the truth of the claims the holder presents. This is not entirely true. To be perfectly precise the holder can prove that the issuer has signed the validity of the claim. For example, a holder can prove that VVV confirms that the subject's name is John Papadopoulos. Whether this claim is true or not, and more importantly whether it can be trusted or not, depends upon the level of trust that the verifier has towards the issuer.

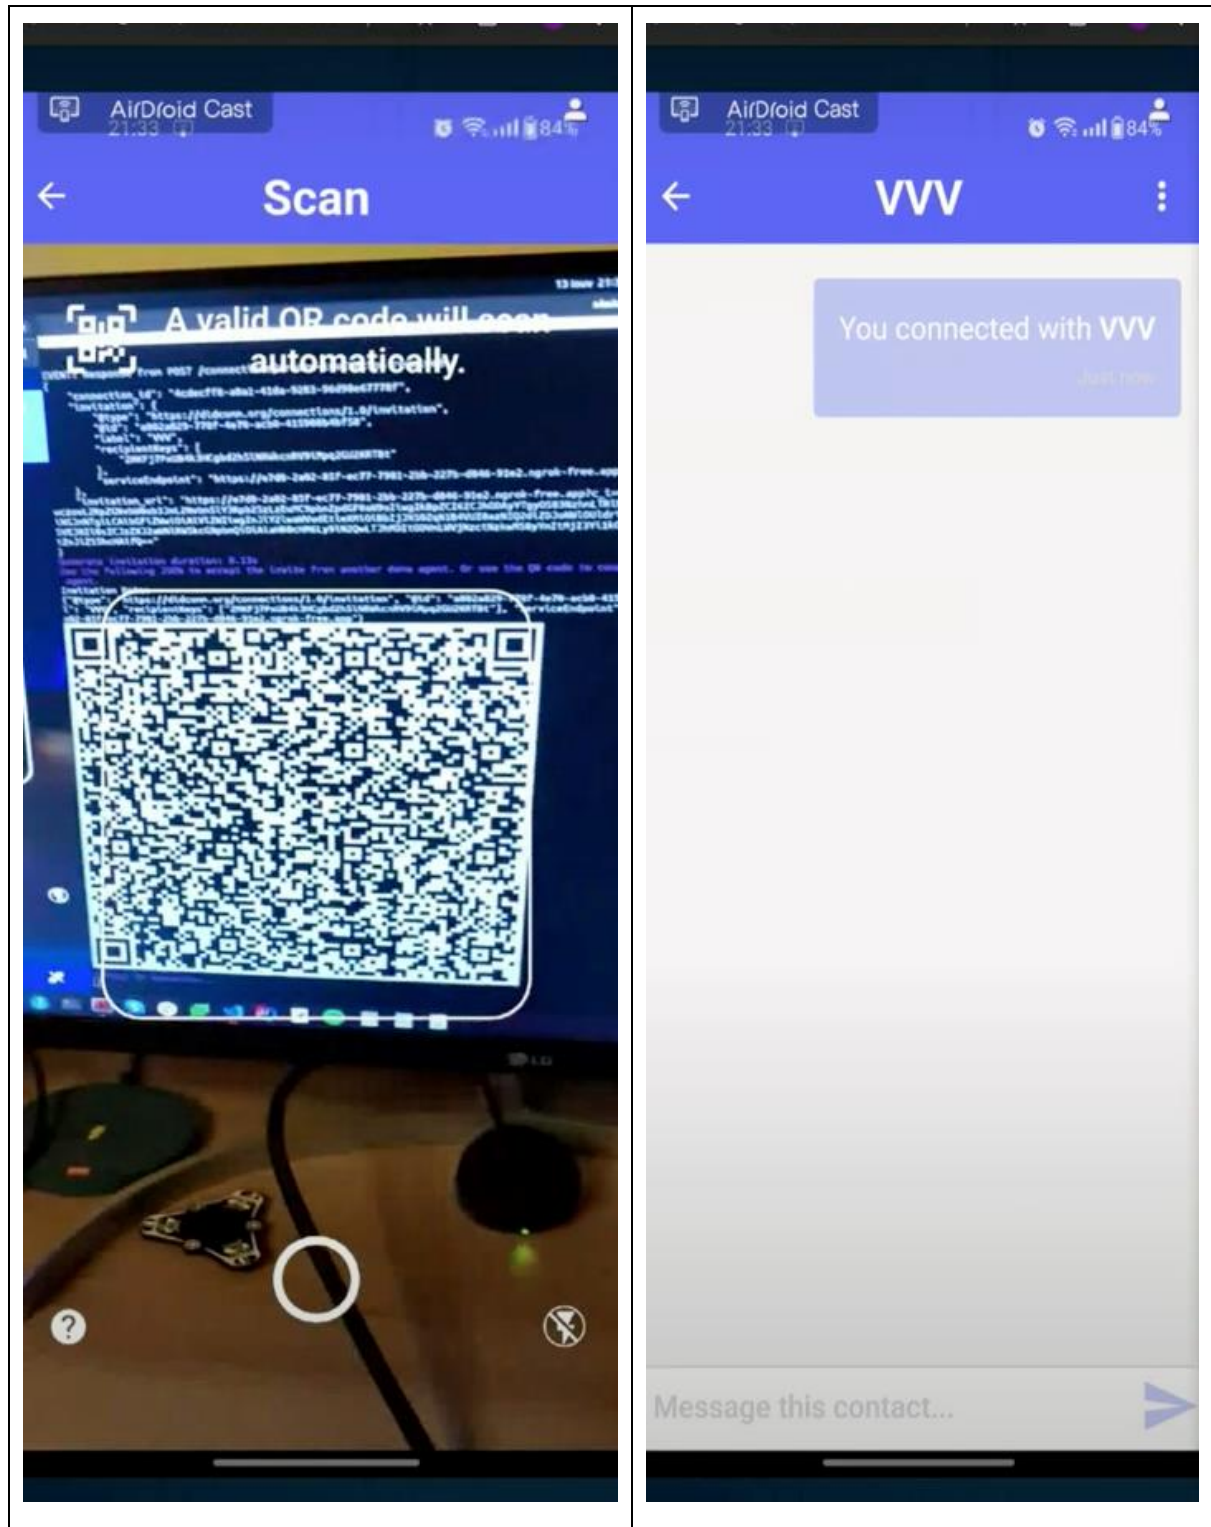


Figure 21: Accepting the invitation

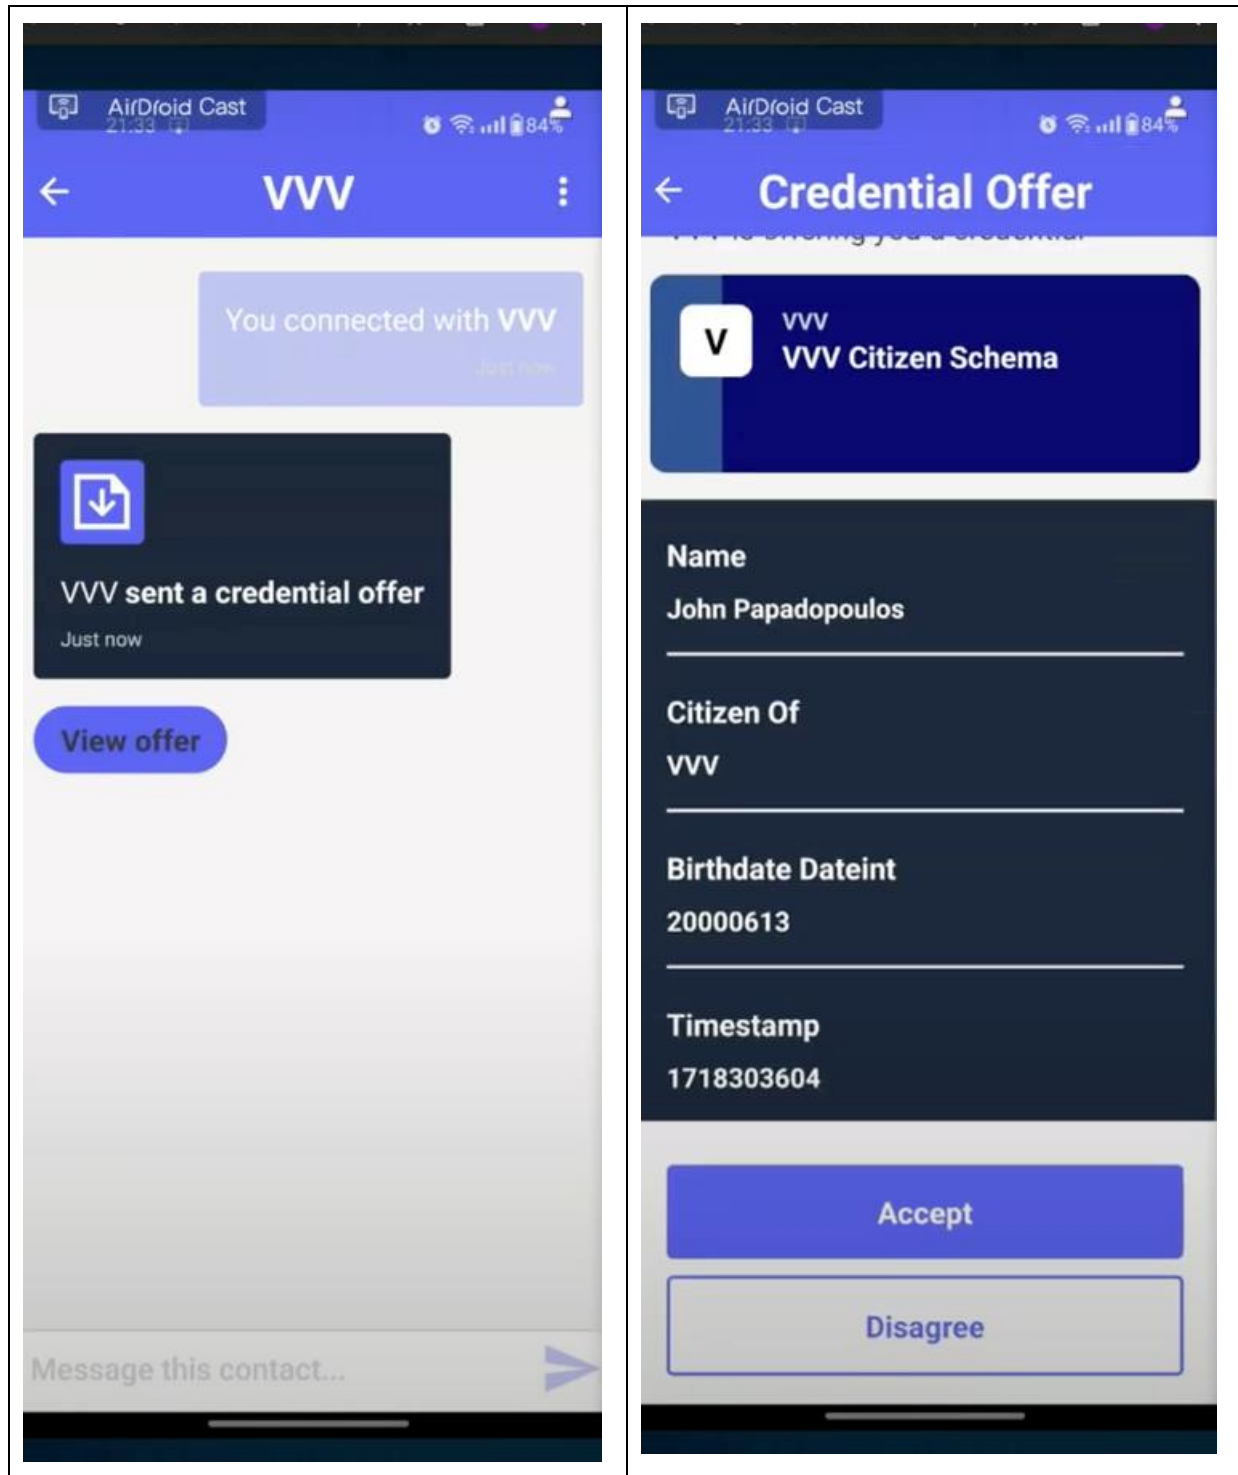


Figure 22: Accepting a credential offer

5 Data Governance Framework

The current section will document the second and final iteration of the Data Governance Framework. The reader can consult Section 5 of D3.1 for the first version of the DGF.

5.1 General Guidelines and Policies

The Data Governance Framework (DGF) is a structured and comprehensive set of guidelines, policies, and procedures that govern how data is managed, shared, and protected within the AI4Gov Project. This framework serves as a crucial instrument for ensuring that data-related activities align with the EU's legal and regulatory landscape, particularly with regard to data protection and privacy. Within this context, the Data Governance Framework project plays a pivotal role in navigating the complexities of data management while complying with EU data protection laws. This framework acts as a structured roadmap that not only empowers project partners to harness the potential of data but also safeguards the rights and interests of individuals whose data is processed.

The DGF is aligned with the Data Governance Act while also taking into consideration key regulations such as GDPR, AI Regulation, EU AI Act and the Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment. The new version of the DGF included in this iteration introduces a comprehensive set of rules and guidelines sourced from the AI Act. This addition strengthens the framework's foundation and expands its scope to cover the specific requirements of AI governance. By incorporating principles from the AI Act, the DGF demonstrates a commitment to high standards of ethical and legal integrity in data handling and AI system deployment within AI4Gov.

To provide a concrete framework, policies and guidelines are generated for each of the above regulations that all partners within AI4Gov should take into consideration, focusing on all of the following factors.

1. **Compliance with Regulations:**

This factor emphasizes the need to comply with data protection and privacy regulations. It includes ensuring that data handling practices align with the legal requirements imposed by such regulations, with a focus on key laws like the General Data Protection Regulation (GDPR).

- Stay informed about relevant data protection regulations in your region and industry.
- Appoint a Data Protection Officer (DPO) to oversee compliance.
- Regularly update data governance policies to align with evolving regulations.

2. **Data Ownership:**

Data ownership refers to the clear definition of who has control over the data. It ensures that data rights and responsibilities are well-defined among partners, especially in collaborative

initiatives. Clear data ownership definitions help prevent disputes and maintain responsible data management.

- Define data ownership in partnership agreements, specifying rights and responsibilities.
- Establish data governance committees with representatives from each partner to address ownership concerns.
- Create data access and usage policies that respect data ownership and provide guidelines for shared data.

3. Data Security:

Data security is vital to safeguard data against unauthorised access and breaches. It involves implementing security measures such as encryption for data at rest and in transit, access controls, and regular security audits to identify and mitigate potential risks.

- Implement encryption for data at rest and in transit.
- Enforce access controls, ensuring that only authorised personnel can access and modify data.
- Regularly conduct security audits and vulnerability assessments to identify and mitigate risks.

4. Data Quality:

Data quality ensures that data used for analysis and decision-making is accurate, consistent, and reliable. It involves the development of data quality standards, validation processes, data profiling, and data cleaning to maintain high-quality data.

- Develop data quality standards and validation processes to maintain data accuracy and consistency.
- Implement data profiling and cleaning procedures to rectify inaccuracies and inconsistencies.
- Provide training to ensure that personnel understand the importance of data quality and their role in maintaining it:

5. Privacy by Design:

Privacy by design emphasises the integration of privacy safeguards into AI development and data handling processes from the project's outset. It includes practices like privacy impact assessments (PIAs) and data anonymisation to protect individual identities.

- Incorporate privacy impact assessments (PIAs) into the development of new projects and data initiatives.
- Use data anonymisation or pseudonymisation techniques to protect individual identities.
- Continuously assess and update privacy measures to adapt to changing risks and challenges.

6. Data Sharing Agreements:

Data sharing agreements are essential for defining the terms and conditions of data sharing, access, and usage. They ensure clarity and compliance in data sharing practices, including specifying data ownership and responsibilities. Regular review and updates are necessary to adapt to changing conditions.

- Draft comprehensive data sharing agreements that clearly specify data ownership, permitted uses, and responsibilities.
- Include provisions for data retention and disposal to maintain compliance with regulations.
- Regularly review and update data-sharing agreements to reflect changing needs and conditions.

7. Data Lifecycle Management:

Data lifecycle management involves a structured approach to data handling, ensuring data consistency from acquisition to disposal. It includes the development of data lifecycle plans, regular audits, and documentation of data retention and disposal processes.

- Develop a data lifecycle management plan to ensure data is handled consistently from acquisition to disposal.
- Regularly audit data storage and processing practices to identify inefficiencies or compliance issues.
- Document data retention and disposal processes to maintain transparency and compliance.

8. Ethical Considerations:

Ethical considerations focus on responsible AI practices and the prevention of bias and discrimination in AI applications. This involves conducting fairness and bias assessments, providing training to raise ethical awareness, and promoting transparency in AI development and deployment.

- Conduct fairness and bias assessments on AI models to identify and mitigate potential bias.
- Provide training to personnel involved in AI and data projects to raise awareness of ethical concerns.
- Encourage transparency by documenting AI model development and deployment processes.

- Take Assessment List for Trustworthy Artificial Intelligence (ALTAI)²⁷ for self-assessment into consideration.

9. **Accountability:**

Accountability ensures clear lines of responsibility within the partnership. It involves the appointment of Data Stewards to oversee data governance, defining roles for incident response and GDPR compliance, and establishing a Data Governance Committee for oversight.

- Appoint Data Stewards within each partner organization to oversee data governance practices.
- Create clear lines of responsibility for incident response and GDPR compliance.
- Establish a Data Governance Committee with representatives from all partners to oversee accountability.

10. **Monitoring and Compliance:**

Continuous monitoring, audits, and compliance assessments are essential to identify and rectify issues, ensuring ongoing data governance. This factor involves regular internal audits, documentation of practices, and the establishment of mechanisms for reporting and addressing data governance concerns and breaches promptly.

- Regularly conduct internal audits and compliance assessments to identify and rectify issues.
- Provide a mechanism for stakeholders to report data governance concerns or breaches for quick resolution.
- Take AI4Gov's Data Management Plan (D1.2) into consideration regarding monitoring activities and compliance.

5.2 Applicable Regulations and EU Guidelines

5.2.1 General Data Protection Regulation (GDPR)

The General Data Protection Regulation (GDPR) is a comprehensive data protection and privacy regulation that was enacted by the EU, replacing the Data Protection Directive (95/46EC). The project's data management plan includes all actions and guidelines for GDPR compliance; since it applies to all data, including decentralised data, it also applies to the Decentralized Data

²⁷ European Commission, Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment, <https://digital-strategy.ec.europa.eu/en/library/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>

Governance Model. With decentralised storage, however, there is a caveat that needs special consideration.

The immutable nature of the blockchain means that any information stored there will always be there and will never be removed as long as at least one node continues to operate and keep a copy of the chain. This seems to contradict the “right to be forgotten” right of GDPR. Even if we assume that personal data is encrypted in the blockchain, there is still the possibility that a future data breach or exploit will break the encryption.

The following considerations and actions are specific to the Decentralized Data Governance Model and aim at fully respecting the “right to be forgotten” tenet of GDPR.

- All files are stored off-chain and can be accessed via the blockchain only through anchors.
- Unencrypted files in plain text will only contain public data (such as publicly available reports).
- Any file that has potentially sensitive information will be anonymised before being stored and anchored; even then, it will be in encrypted format.
- Anonymised and encrypted files can only be unencrypted by users owning the appropriate key pair. Only the data controllers will have such keys.
- Anonymised and encrypted files can be deleted by the data controllers from the off-chain storage. The anchor will become invalid and will not be able to get verified.
- Distributed storage technologies, such as IPFS, will only be used for public data.

On top of the above points, the OpenDSU mechanism of the wallet performs further encryption and sharing of data and allows fine-tuned data control to the wallet’s owner, thus applying further data protection for end-users.

GDPR Compliance Guidelines

1. Legal and Regulatory Compliance (Article 5):
 - Ensure that all data processing activities comply with the GDPR and relevant data protection regulations.
2. Data Classification and Sensitivity (Article 5):
 - Classify data based on sensitivity and importance to determine appropriate safeguards and handling requirements.
3. Data Protection Officer (DPO) (Articles 37 & 38):
 - Appoint a Data Protection Officer if required by the GDPR and define their responsibilities.
4. Data Inventory and Mapping (Article 30):
 - Create a comprehensive data inventory and mapping to understand data flows, storage locations, and processing purposes.
5. Data Minimization (Article 5):
 - Collect and process only the data necessary for the project's objectives, adhering to the principle of data minimization.
6. Data Subject Rights (Articles 12-23):
 - Ensure that data subjects can exercise their rights, such as the right to access, rectify, and delete their data.
7. Data Processing Legal Basis (Articles 6 & 9):

- Identify and document the legal basis for data processing activities within the project.
- 8. Data Protection Impact Assessments (DPIAs) (Article 35):
 - Conduct DPIAs for high-risk data processing activities and take measures to mitigate identified risks.
- 9. Privacy by Design and by Default (Article 25):
 - Integrate privacy into the project's design and development processes to ensure data protection is a core consideration.
- 10. Data Security (Article 32):
 - Implement strong data security measures, including encryption, access controls, and regular security audits.
- 11. Data Breach Response Plan (Articles 33 & 34):
 - Develop a clear and documented plan for responding to and reporting data breaches in compliance with the GDPR.
- 12. Consent Management (Article 7):
 - If applicable, establish a consent management system for collecting, recording, and managing consent from data subjects.
- 13. Third-Party Data Processors (Article 28):
 - Ensure that any third-party data processors involved in the project comply with GDPR and have appropriate data processing agreements in place.
- 14. Data Transfer Mechanisms (Chapter V):
 - Implement lawful mechanisms for international data transfers, such as Standard Contractual Clauses (SCCs).
- 15. Data Retention and Deletion (Article 5):
 - Define data retention policies and procedures to ensure data is not kept longer than necessary for the intended purposes.
- 16. Data Access and Portability (Article 20):
 - Provide mechanisms for data subjects to access and receive their data, adhering to GDPR's data portability requirements.
- 17. Training and Awareness (Article 39):
 - Conduct training for project stakeholders to increase awareness of data protection principles and GDPR compliance.
- 18. Data Governance Policies and Procedures (Article 5):
 - Develop clear data governance policies and procedures that outline how data should be managed and processed within the project.
- 19. Data Sharing Agreements (Article 28):
 - If data sharing occurs with external entities, establish clear data sharing agreements that include data protection clauses.
- 20. Data Documentation and Records (Article 30):
 - Maintain detailed records of data processing activities, agreements, and compliance measures.
- 21. Regular Auditing and Monitoring (Article 32):
 - Implement regular audits and monitoring to ensure ongoing compliance with the GDPR and other data protection regulations.

22. Incident Response Plan (Articles 33 & 34):

- Develop a response plan for handling and reporting data incidents as required by the GDPR.

5.2.2 EBSI Conformance

Created in 2018, the European Blockchain Services Infrastructure (EBSI) is the EU's "official" blockchain infrastructure. It operates with nodes across EU countries with the goal of offering its services to organisations and citizens across Europe. Its business use cases currently aim at three domains, namely Verifiable Credentials, Track and Trace and Trusted Data Exchange. During the EBSI demo day, held in May 2022, various scenarios proved the ability to verify credentials using the underlying EBSI infrastructure; that means that EBSI, though still an active and ever-growing project, has proved its efficiency for cross-border credential certification.

If the trend continues, EBSI will be adopted in production, and it will be the main infrastructure for cross-border, SSI-enabled, cross-border transactions. As such, potential integration with the AI4Gov blockchain infrastructure and dApp ecosystem will be investigated during the design and implementation of the smart contracts and dApps required for the execution of the pilot use cases.

Two main aspects, however, can be identified at the present moment:

- **Wallet conformance.** A goal that can be set from the present moment is that the wallet that will be implemented for AI4Gov will conform to EBSI standards. This conformance is verified by a series of tests offered by EBSI. The tests differ depending on the role of the Wallet user (end-user or holder, issuer, verifier). For AI4Gov, it is expected that Holder Wallets will be implemented; however, if any issuer or verifier wallet application is needed, this, too, shall be tested for EBSI conformance. All DIDs used in AI4Gov will be fully compliant with the EBSI guidelines.
- **Usage of EBSI services.** This is related to the first one in the sense that EBSI services can be used only by conformant applications. This aspect will investigate which of the EBSI services that are offered or planned to be implemented can be used for AI4Gov (Identity provision via an authorisation endpoint or via the SSI eIDAS bridge is an example).

5.2.3 Ethics Guidelines for Trustworthy AI

The Ethics Guidelines for Trustworthy Artificial Intelligence was presented in April 2019 by EU's High-Level Expert Group on AI.²⁸ The Guidelines put forward a set of 7 key requirements that AI systems should meet in order to be deemed trustworthy.

²⁸ European Commission, Ethics guidelines for trustworthy AI, 2019, <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>

1. Human agency and oversight:
 - AI systems should empower human beings, allowing them to make informed decisions and fostering their fundamental rights. At the same time, proper oversight mechanisms need to be ensured, which can be achieved through human-in-the-loop, human-on-the-loop, and human-in-command approaches.
2. Technical Robustness and safety:
 - AI systems need to be resilient and secure. They need to be safe, ensuring a fall back plan in case something goes wrong, as well as being accurate, reliable and reproducible. That is the only way to ensure that also unintentional harm can be minimized and prevented.
3. Privacy and data governance:
 - Besides ensuring full respect for privacy and data protection, adequate data governance mechanisms must also be ensured, taking into account the quality and integrity of the data and ensuring legitimised access to data.
4. Transparency:
 - Data, system and AI business models should be transparent. Traceability mechanisms can help achieve this. Moreover, AI systems and their decisions should be explained in a manner adapted to the stakeholders concerned. Humans need to be aware that they are interacting with an AI system and must be informed of the system's capabilities and limitations.
5. Diversity, non-discrimination and fairness:
 - Unfair bias must be avoided, as it could have multiple negative implications, from the marginalisation of vulnerable groups, to the exacerbation of prejudice and discrimination. Fostering diversity, AI systems should be accessible to all, regardless of any disability, and involve relevant stakeholders throughout their entire life circle.
6. Societal and environmental well-being:
 - AI systems should benefit all human beings, including future generations. It must hence be ensured that they are sustainable and environmentally friendly. Moreover, they should take into account the environment, including other living beings, and their social and societal impact should be carefully considered.
7. Accountability:
 - Mechanisms should be put in place to ensure responsibility and accountability for AI systems and their outcomes. Auditability, which enables the assessment of algorithms, data and design processes plays a key role therein, especially in critical applications. Moreover, adequate and accessible redress should be ensured.

In the scope of the Data Governance Framework, a questionnaire has been created based on the Assessment List for Trustworthy Artificial Intelligence (ALTAI) for self-assessment; please refer to Appendix C of D3.1 for the questionnaire.

5.2.4 EU Artificial Intelligence Act

The EU's Artificial Intelligence Act (AI Act) was formally passed by the European Parliament on March 13, 2024, and received final endorsement from the European Council on May 21, 2024. The AI Act will enter into force 20 days after its publication in the Official Journal of the European Union, which is expected in June or July 2024. The AI Act introduces a risk-based framework to regulate AI systems, categorizing them into unacceptable, high, limited, and minimal risk levels, each with specific compliance obligations. High-risk AI systems will be subject to stringent requirements, including conformity assessments and detailed documentation.

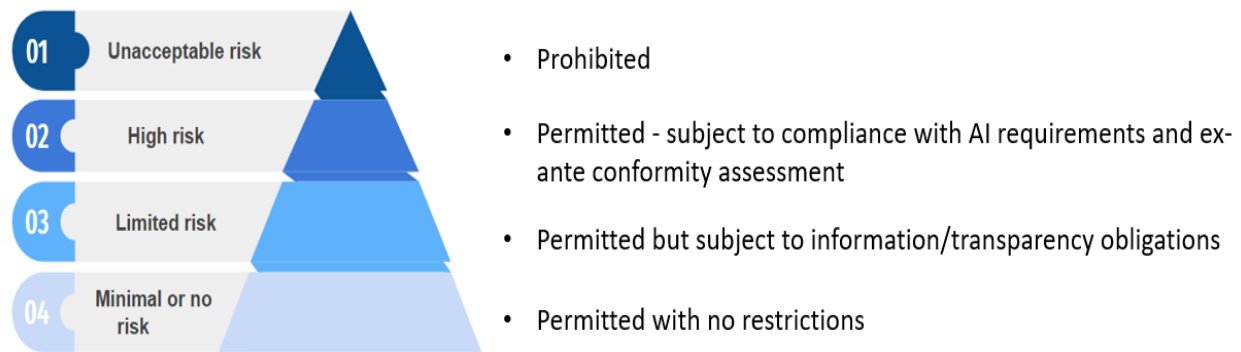


Figure 23: AI Act defined levels of risk

First proposed in April 2021, the European Commission proposed the first EU regulatory framework for AI as part of the EU's broader efforts to address the ethical and legal challenges posed by AI technologies. The AI Act serves as a comprehensive regulatory framework within the European Union, designed to standardize the development and deployment of artificial intelligence technologies. Its central mission is to ensure that AI systems are created and employed in a manner consistent with the ethical, legal, and safety standards upheld by the EU. The AI Act encompasses a wide range of AI applications, classifying them as either high-risk or low-risk based on their potential to cause harm. The regulation expressly prohibits certain AI practices, including government-based social scoring and the exploitation of individuals' vulnerabilities.

Transparency and accountability are also foundational principles of the AI Act, necessitating clear documentation, user information, and explanations for AI system decisions. Data governance, focusing on data quality and data protection, is also a critical component of the regulation. For non-compliance with its provisions, the AI Act outlines penalties and fines, with monetary penalties that can range up to €30 million or 6% of the violating entity's global annual turnover, contingent on the severity of the breach.

A brief analysis of the AI Act articles content which are incorporated in the DFG in this second iteration is presented below:

Article 1-3: Scope and Definitions - *Defines the scope, objectives, and key definitions of AI systems covered by the Act.*

Article 4-5: General Provisions - *Establishes the foundational principles for AI system compliance, including risk management and ethical considerations.*

Articles 6-8: Risk Classification - *Details the classification of AI systems into risk categories and associated requirements.*

Articles 9-11: Data Management - *Outlines requirements for data quality, protection, and management in AI systems.*

Articles 12-14: User Rights and Transparency - *Specifies user rights and obligations for transparency in AI system operations.*

Articles 15-17: Human Oversight and Security - *Emphasizes the need for human oversight and robust security measures.*

Articles 18-20: Impact Assessments and Governance - *Details the requirements for conducting impact assessments and establishing governance frameworks.*

Articles 21-23: Incident Response and Training - *Provides guidelines for incident response, reporting, and stakeholder training.*

Based on the analysis of the articles above, the DGF is extended with the following principles and guidelines for each relevant field, all of which should be taken into account by technical partners introducing new AI apps and solutions with the Ai4Gov project.

1. Legal and Regulatory Compliance (Articles 1 & 5)

- Ensure all AI systems developed within the AI4Gov project comply with the AI Act
- Regularly update compliance procedures to reflect changes in recent AI legislation

2. Risk Management and Classification (Articles 6,7 & 8)

- Classify AI systems based on their risk level (unacceptable, high, limited, minimal).
- Implement appropriate safeguards for each risk level.

3. Transparency and Documentation (Articles 13 & 14)

- Maintain comprehensive documentation for all AI systems, including design, purpose, and compliance measures.
- Ensure transparency by providing clear information to users and partners about AI system operations and decisions.

4. Data Protection and Privacy (Articles 10 & 11)

- Implement data protection measures to align with GDPR
- Ensure personal data used by AI systems is anonymized or pseudonymized where possible.

5. Human Oversight and Accountability (Articles 14 & 15)

- Establish mechanisms for human oversight to monitor AI system performance and decisions with AI4Gov
- Assign accountability for AI system operations and ensure human intervention is possible.

6. Data Quality and Management (Articles 8 & 9)

- Ensure high-quality data for training, validation, and testing AI systems
- Implement data management practices to maintain data integrity and accuracy

7. Robustness and Security (Articles 16 & 17)

- Develop AI systems with robust security measures to protect against cyber threats and vulnerabilities
- Regularly test and update security protocols set by the partner organization to ensure resilience

8. Impact Assessments and Mitigation (Article 18)

- Conduct impact assessments for high-risk AI systems to identify and mitigate potential risks
- Develop and implement risk mitigation plans based on assessment findings, informing involved partners

9. User Rights and Consent (Article 12)

- Ensure AI system users can exercise their rights, such as access, correction, and deletion of data.
- Obtain explicit consent from users for data processing if required within the project

10. Ethical Considerations (Article 4 & 5)

- Avoid AI applications that could result in discrimination, bias, or unfair treatment.
- Integrate ethical principles into AI system design and operation.

11. AI Governance and Monitoring (Articles 19 & 20)

- Establish governance structures to oversee AI system development and deployment.
- Implement continuous monitoring and evaluation processes to ensure compliance and performance.

12. Incident Response and Reporting (Articles 21 & 22):

- Develop incident response plans to address and report AI system failures or breaches.
- Ensure timely reporting to project coordinator, relevant authorities and affected parties.

13. Training and Awareness (Article 23)

- Provide training for partner stakeholders on AI Act compliance and ethical AI practices.
- Promote awareness of AI system risks and benefits among users and developers within the AI4Gov ecosystem

6 Conclusions

The present report provided the final iteration of the Decentralised Data Governance framework for AI4Gov. All the major technology components that will facilitate decentralization in the pilot use cases have been implemented. A major addition was the possibility to allow for citizens to join the platform and leverage the blockchain capabilities to participate in open democracy schemes. The general mechanism that allows such self-governing bodies to form and exist is the so called Digital Autonomous Organization or DAO. The architecture and the prototype components have been expanded to allow for this new functionality.

It is to be noted that the term DAO is an umbrella term similar to the term smart contract. It does not refer to any specific functionality but to components implemented via smart contracts and can be self-governed. As such, the major addition to the framework was the capability to host such DAOs. This capability was demonstrated by instantiating a specific DAO that allows citizens to propose new policies and vote on these new policies. However, custom DAOs can be implemented to realize different use case scenarios.

Under these considerations, the architecture was redesigned to accommodate the new functionality, while citizen-centric wallets using the HyperLedger Aries framework have been implemented. The OpenDSU technology, which has been considered initially for the wallet implementation, was substituted with the new approach due to reasons having to do both with the measured performance of the OpenDSU-based prototypes and for having a common implementation scheme for both types of wallets, organizational and citizen.

Lastly, the second and final iteration of the Data Governance Framework was consolidated and presented in the present report.

7 References

Liu, J., Makarov, I., & Schoar, A. (2023). *Anatomy of a run: The terra luna crash*.
<https://www.nber.org/papers/w31160>