Centre for the Fourth Industrial Revolution

WORLD ECONOMIC FORUM

# Bridging the Governance Gap: Interoperability for blockchain and legacy systems

**WHITE PAPER**

DECEMBER 2020

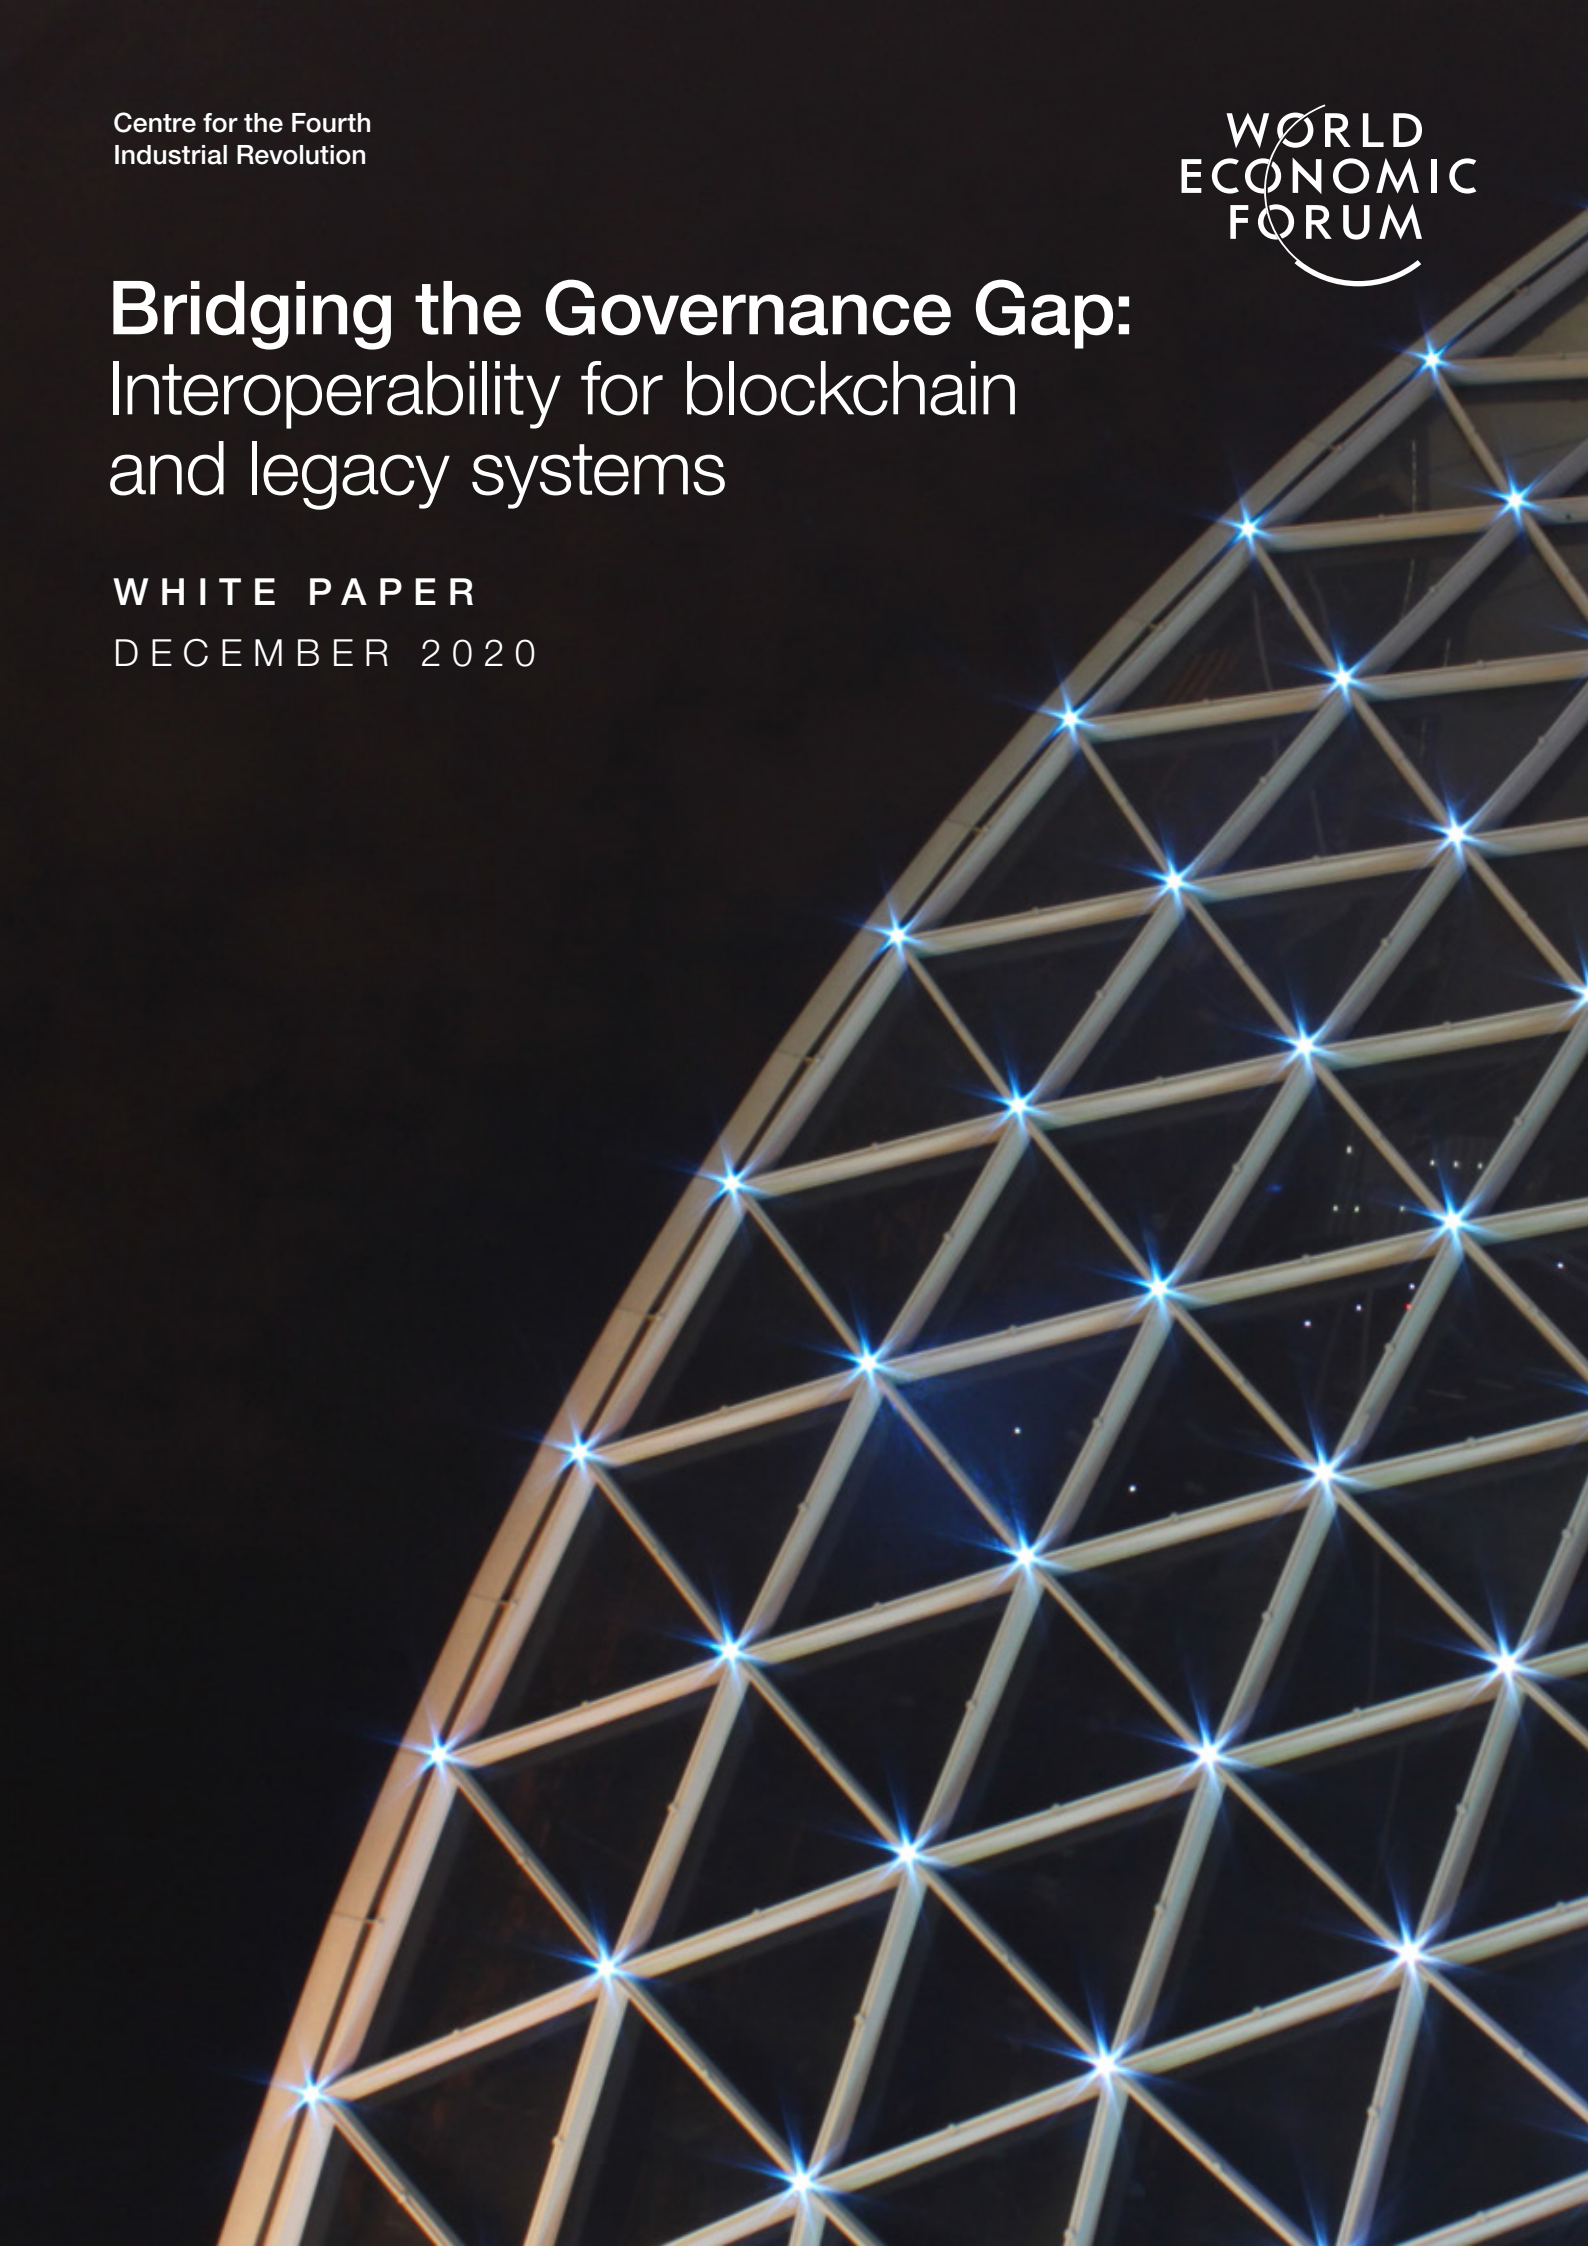

**Cover:** Unsplash/Terry

**Inside:** Unsplash/chuttersnap; Unsplash/Luke Ellis; Unsplash/RenRan; Unsplash/MarkusSpiske; Unsplash/ScottWebb; Unsplash/AndersJilden; Unsplash/ChristianHolzinger

# Contents

# 1 Introduction

Smart contracts can unlock the hidden value of legacy digital systems based on interoperability with the capabilities of DLT systems.

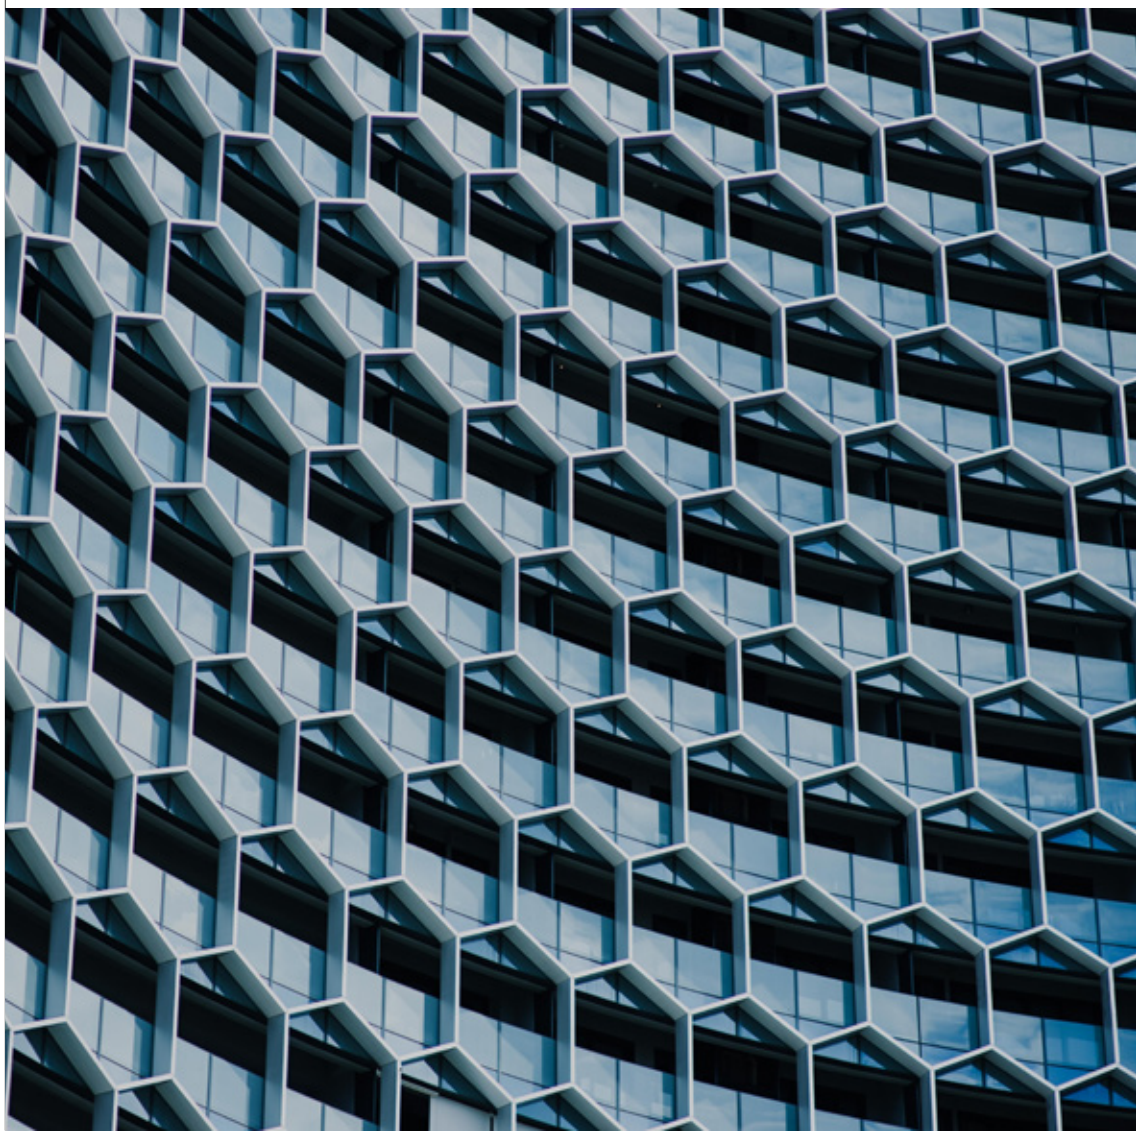

In the past 10 years, blockchain and distributed ledger technologies (DLT) have generated tremendous interest and activity from developers, enterprises, venture capitalists, regulators and users alike. The innovation of blockchain technology leads to an important question about how legacy digital systems, operated by enterprises, governments and institutions, will be affected. Presently, the answers to this question have varied from one extreme ("all legacy systems will be replaced") to the other ("DLT is too slow and unproven to actually replace any working legacy system"). However, the eventual answer may lie somewhere in between, where the utility of select legacy systems is upgraded by DLT integration wherever appropriate, and DLT solutions witness a growth in enterprise adoption.[1]

Multiple reports analysing the blockchain/DLT adoption by organizations have pointed out that blockchain integration with other systems (e.g. other blockchains or other non-DLT information systems) is one of the crucial challenges. Early experiments on interoperability have demonstrated DLT-to-legacy integration to be useful in many use cases (e.g. reinsurance) in establishing trust among multiple parties. This white paper intends to identify the foundational pillars for legacy system-DLT system interoperability as well as efforts made in this direction, and recommends the adoption of a holistic industry standard that can accelerate the development and implementation of these "interoperability bridges" across geographic regions and disparate systems.

## 1.1 Why do legacy systems and DLT solutions need to embrace each other?

Blockchain and distributed ledger technologies (DLTs) are backend infrastructure that store and transfer data among parties within a shared ledger without the need for traditional intermediaries. Since the ledger is redundantly processed and updated by a decentralized network of computers instead of a centrally managed server, users can generally trust the integrity of the data and computations without the need for trusted authorities. The objective is to use a blockchain as a shared ledger for exchanging value between disparate entities in order to achieve greater efficiency, heightened transparency and a reduction in complex reconciliation.

One of the most discussed innovations in blockchain technology is smart contracts – conditional business logic (if *x* event happens, then execute *y* action) that is executed on the blockchain. These also allow the creation of digital assets where the smart contract uses embedded logic by defining how those digital tokens function (e.g. represent voting rights or a stake in protocol revenue). Since smart contracts operate on a blockchain, in general no counterparty or external entity can tamper with the terms, execution or outcome, providing the user with technologically enforced guarantees that the contract will be fairly honoured.[2] Given the nature of these guarantees, smart contracts are being looked at as a new form of multiparty business automation. However, it should be noted that smart contracts are vulnerable to exploitation if their code is not audited well for security, fault tolerance (even when nodes go down) and privacy.

As part of their underlying design, blockchains create strong security guarantees at the expense of being inherently disconnected from any network or system in the outside (off-chain) world. This decentralized consensus design runs on other inherent costs such as rewards for validating transactions on-chain and ensuring fault tolerance etc. It has been referred to as one of the features of "blockchain trilemma": ensuring blockchain decentralization, scalability and security imply trade-offs, at least in the short term. To a certain extent, the mechanisms for ensuring decentralization at different blockchain layers may conflict with security and scalability.

While simple smart contracts can be created within isolated blockchain environments to store data or execute transactions between users, they stand to provide more value to enterprise and other ecosystems when connected to data and systems outside of the blockchain. For instance, if an insurance service provider wishes to automate the dispersal of flight insurance claims, it could do so via a smart contract. But for the smart contract to execute, it needs accurate information on the scheduled and actual departure times of the flights, as well as the ability to settle in fiat currencies on traditional payment rails. Since blockchains do not natively generate this information or provide connections to traditional systems, a piece of infrastructure called an "oracle" must be adopted to connect the smart contract to external resources. Additionally, since a smart contract cannot execute on software outside of the blockchain, the oracles provide the smart contract with inputs and also take the outputs and execute them as actions on external systems.

Oracles serve as an external source of data for DLT systems. They connect the external world to the self-contained world of DLTs, acting as middleware for data and transaction sharing between environments in a secure and authoritative manner. Information shared by oracles is digitally signed and hence is considered non-repudiable (assurance that the signature cannot be denied by the party who signed it). Since smart contracts are executed on DLT in a deterministic fashion (i.e. transactions

that execute exactly as written because they are verified by all nodes with high fault tolerance), the built-in ability to communicate with the external world becomes difficult to establish without a loss of trust. Oracles become this entity that signs claims about the state of the external world. Since DLT systems were built intentionally to be detached from the external world and its trusted third parties, it is crucial that the links (i.e. oracles) have high integrity. Decentralization, economic incentives, use of trusted execution environments etc. are some of the ways in which oracle integrity is ensured on a varying basis (based on the end use case). Some of these approaches are discussed in this paper.

Oracles are not blockchains themselves but are **secure blockchain middleware** that operates partly on the blockchain ("on-chain") and partly outside of the blockchain ("off-chain") asynchronously. The main goal of an oracle is to retrieve external data, validate it and deliver it to the intended entity. Each step of the process requires important considerations to ensure the desired qualities of the blockchain (e.g. being tamper-resistant, permissionless and immutable) are not lost when expanding to off-chain data and systems. This requires an oracle to be able to source data from high-quality application programming interfaces (APIs), show proof of the origin of the data, maintain secure and reliable data delivery, provide economic guarantees to incentivize trusted oracle services and possibly even provide additional cryptographic techniques to keep the data itself private. Some oracles are physical or tangible devices that measure real-world values, such as temperature or whether a shipment has arrived safely. Other oracles are intangible, and comprise only code. For instance, the recording of external product prices on a blockchain is facilitated by oracles.

FIGURE 1 **Data flow in a typical DLT-legacy interoperability framework through oracles**

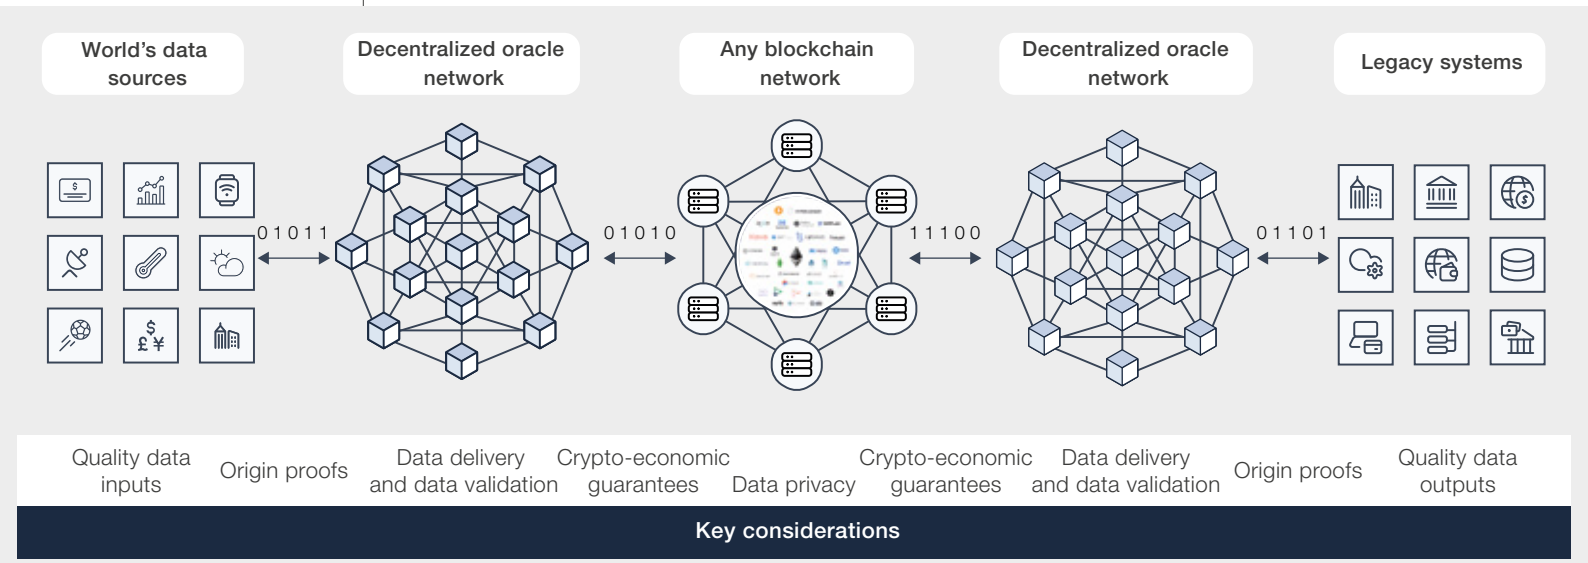


An enterprise's legacy systems could generally connect, bidirectionally, to blockchain networks using an oracle. This model is essential if smart contracts are to have a significant global impact on business process efficiency and transparency, particularly because enterprise smart-contract use cases generally require access to high-quality off-chain data and traditional business infrastructure in order to run end-to-end.

In order to ensure communication between off-chain legacy systems and on-chain smart contracts and to be able to expand the utility of smart contracts beyond just DLT applications, adopting an approach towards building secure middleware solutions (developed using open-source technology platforms) that anyone is able to connect to and build on as needed, along with governance models that harmonize inter-system communication, is required.

Let's explore an example to understand why such interoperability is beneficial.

The Government of India launched an ambitious programme called Pradhan Mantri Fasal Bima Yojana (translated as Prime Minister's Crop Insurance Scheme; abbreviated as PMFBY) in 2016 to provide farmers with insurance coverage and financial support for any losses due to unforeseen crop failures arising out of natural calamities/pests/diseases and to encourage them to adopt modern technology and innovative farming methods. Most of the insurance premium is covered by the central and state governments, while the farmer pays a maximum of 2% of the total premium annually. The large scale of the programme is evident from its broad coverage and the type of stakeholders involved, as reported by the government in the Kharif cropping season (June–October) of 2019.[3]

FIGURE 2 **Highlights of Pradhan Mantri Fasal Bima Yojana**

| | |
|---|---|
| Total farmers covered | 18 million |
| Crop area insured | 28 million hectares |
| Total premium | $2 billion |
| Total insurance companies involved | 12+ |

**Stakeholders / implementational agencies**

| | | |
|---|---|---|
| Central govt. | Insurance cos. | Weather cos. |
| State govt. | Banks | Farmers |
| Insurance regulator | Agri research institutions | Drone / remote sensing cos. |

**Source:** The Prime Minister's Crop Insurance Scheme as on October 2019

In accordance with the operational guidelines, a variety of external stakeholders are brought together to generate/collect data on crop yield and weather, conduct crop-cutting experiments to verify actual yield, undertake remote sensing to collect large-scale data on crop impact, working with insurance companies to underwrite the risk, and banks/governments to provide farmers with credit and welfare transfers. All of these stakeholders and many other subcontracted agencies operate a variety of technical systems to collect and process data, which often have highly varying requirements in terms of privacy, security, transaction speed and human intervention (usually mandated by governmental regulations). Apart from these systems, the central government operates a National Crop Insurance Portal as the central IT system that interacts with all of the other infrastructure and external data sources required. The following schematic shows a summary of the many different types of data sources that are used in the operationalization of PMFBY:

FIGURE 3 **Various legacy IT systems being used in PMFBY and their data**

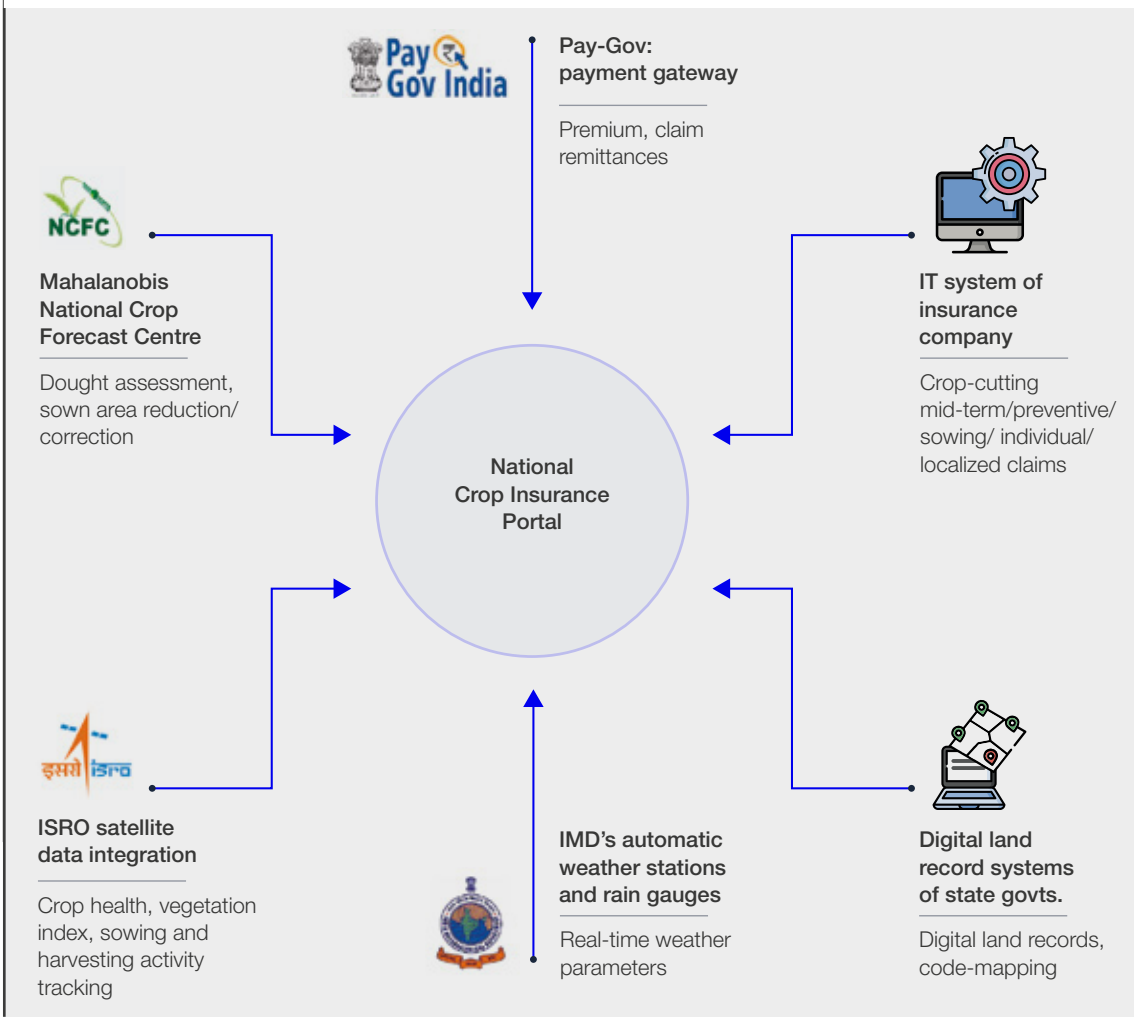


**Source:** https://pmfby.gov.in/pdf/Revised_Operational_Guidelines.pdf

## 1.2 Systemic concerns and how integration with DLT can help

Given the scale of the scheme and the multiplicity of stakeholders involved, several concerns are identified and addressed by the implementational agencies:

- **Transparency and accountability**: As a taxpayer-funded programme, transparency and accountability become increasingly important for the government and regulators. An advantage blockchain brings over legacy databases is its ability to record the first use of data (in a transaction or publication) on the blockchain, record it immutably (so there is no ability to change it later or, at least, if changes are made, they are publicly viewable) and allow open verification by all participants. The "nodes" of the blockchain can be spread among various government (including ombudsman bodies) and civil-society stakeholders in a permissioned blockchain network. Generally, assuming the integrity of data coming from the various agencies is to be trusted, this network can serve as a trustworthy, fault-tolerant, common data store for participants.

- **Reducing corrupt behaviour using smart contracts**: DLT-based smart contracts can be programmed to automatically transfer welfare benefits to farmers based on whether the data relayed by legacy systems and/or stakeholders satisfy preset conditions (e.g. a certain amount of rainfall occurred). This largely eliminates human discretion in settling claims and improves the claim settlement time and cost via data-driven automation (as projected in Figure 1).

- **Data validation through oracle systems**: Since the programme depends on collecting data from multiple independent agencies, concerns about the data's validity and reliability are always present. Using oracles that draw data from multiple sources to process, validate and relay information from external systems to be stored as immutable records on a blockchain can help ensure that key datasets are securely transmitted from authentic sources and validated against tampering before the execution of any agreements occur. However, there still need to be protections for vulnerabilities with regards to data sources being compromised. (This concept is explained in depth later in this white paper.)

- **Information security**: Keeping digital information secure solely through traditional methods may not be very effective, especially against rogue system administrators and single points of failure. DLT generally makes it difficult to break in and alter records (especially on large public networks such as Bitcoin and Ethereum), which may further incentivize stakeholders to provide honest services to maintain a strong reputation.

The crop insurance programme serves as an apt case to highlight the current deficiencies most legacy systems face when dealing with multiparty business processes. While a blockchain should not necessarily be deployed merely for digitizing and coordinating data, if an organization's objective is to automate business processes in a decentralized and disintermediated manner for various reasons, then a blockchain-based architecture becomes imperative. Crop insurance, as depicted above, requires a number of entities to share and verify information (about estimates of crop damage, historical yield, local weather information, estimated loss of production, neighbouring farms and their yield data etc.). In most cases, the interests of these entities are not totally aligned, which is often by design, to ensure that mechanisms for checks and balances work. However, the data relating to the farmer, crop damage and yield history must be agreed upon by all parties as it is integral to the use case. Disintermediation using blockchain allows consensus on data to be reached by all parties (meaning, in this case, claims are settled only when all parties agree on the extent of the crop damage). Additionally, blockchain also provides the capability of maintaining a high level of data integrity (e,.g,. censorship-resistant and tamper-proof record-keeping) and has a built-in high technical fault tolerance. As a result, well-audited smart contracts can infuse the business processes with new capabilities to evaluate insurance claims faster and more accurately, and deliver benefits to the intended recipients more quickly. A report by the Food and Agricultural Organization (FAO) outlines the opportunities of smart contracts and legacy-DLT interoperability in agriculture insurance and micropayments.

FIGURE 4 | **FAO Report: Blockchain for Agriculture (2019)**

**Agricultural insurance** systems in the Asia-Pacific region range from major public-sector programmes in India and the Philippines through to public-private partnerships in China and the Republic of Korea and finally to the purely private markets encountered in Australia and New Zealand and the non-formal private mutual and community-based crop and livestock initiatives in Bangladesh, India and Nepal. Low-cost agricultural insurance schemes are increasingly viewed as mechanisms for providing social protection to the increasing numbers of people affected by floods or droughts and helping to lessen the impacts they suffer as a result of such events. However, despite the multiple benefits, the rate of adoption of insurance products by the rural poor still remains relatively low. The mechanisms that are in place to validate claims and to effect payouts are still time-consuming and this is one of the reasons for index-based insurance not being chosen as the first risk-mitigation strategy by smallholder farmers. Index insurance based on smart contracts can automate and greatly simplify the process, thereby facilitating instant payouts to the insured in the case of adverse weather incidents. Automatic data feeds provide continuous and reliable hyperlocal data to the contract, thereby eliminating the need for on-site claim assessment by the surveyor.

**Source:** http://www.fao.org/3/CA2906EN/ca2906en.pdf

## 1.3 Limitations of DLT and challenges of integration

In addition to the above opportunities and benefits, it is important to consider the limitations of blockchain and DLT and the trade-offs that need to be made.

- **Resource consuming and limited scalability:** A blockchain spends more resources on consensus protocols (computation and participating nodes) to reduce the risk of double-spend attack (more prominent in permissionless public blockchains). Also, on large public blockchain networks, transaction verification takes time,; therefore, throughput (the number of transactions being validated and added to the ledger per second) is currently lower than traditional systems. Technology researchers have been working towards scalability solutions for permissionless blockchains, although it is unclear when these will be achieved. Permissioned blockchain networks generally do not have scalability or resource consumption challenges.

- **Data dissemination risk**: Enterprises are concerned with unintended data dissemination risk on blockchains. Some of these concerns can be partially or fully addressed through private side chains (e.g. Constellation used in Quorum network)[4] or advanced cryptography techniques (e.g. zero-knowledge proofs), but may be costly to implement in the current state. Baseline Protocol also tries to address this concern by using zero-knowledge proofs to allow enterprises to prove that each counterparty used the same common set of records and functions within a shared business process without internal data leaving their respective databases.

- **Unclear regulation**: As discussed in later sections, one of the interoperability framework pillars suggests using crypto-economic guarantees in relation to digital assets. In many jurisdictions, clear regulatory standards are yet to be implemented for the use of such assets, particularly in cross-border transactions.

Blockchain technology is currently evolving, as is the regulatory understanding of its wide-ranging implementations. This paper assumes that readers understand the technology's general capabilities and limitations in detail, and it is not intended to be a framework to assess whether a specific use case is suitable for blockchain. Instead, once readers have already established that blockchain is desirable for their specific use case and business processes, this paper aims to spotlight the role of blockchain, smart contracts and oracles in accelerating the automation of such processes. It highlights how an oracle-powered abstraction layer introduced on top of legacy systems and blockchains can improve the integrity of the interoperability among systems. The recommendations made in the paper thus comprise fundamental pillars in building such an interoperability bridge between legacy and DLT systems, so that all of the systems can be held to the same standard in terms of security and integrity.

# 2 Definitions

Clear definition of blockchain terminology is crucial to understanding and conveying ideas precisely.

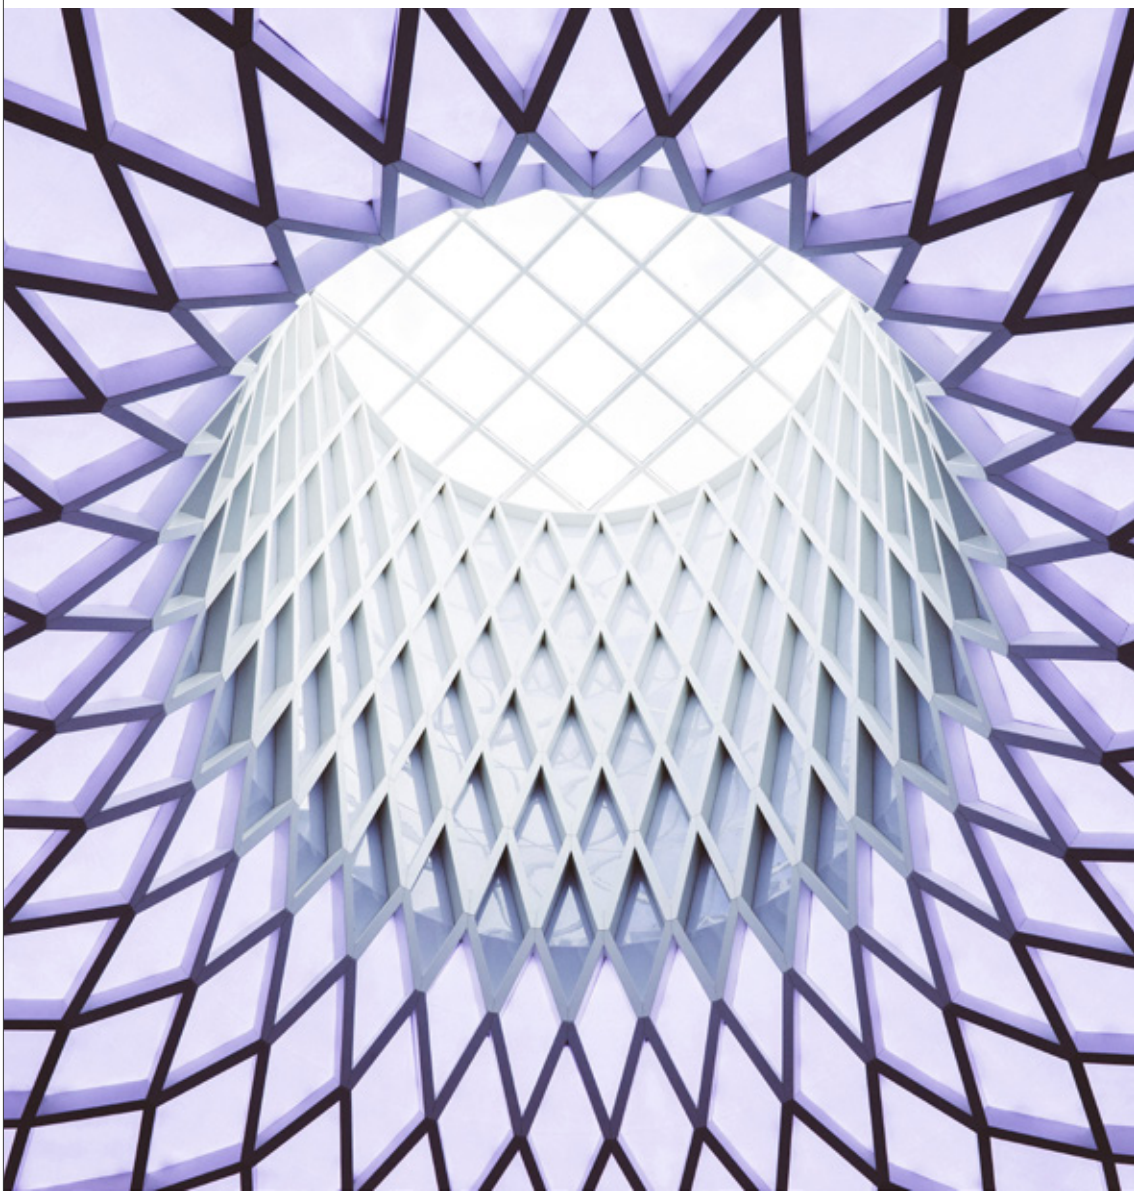

Many of the terms used in this white paper are well understood but tend to have context-dependent interpretations. The following list describes the frequently used terms and articulates the context in which they should be interpreted within the text, unless otherwise indicated.

**Blockchain network**

- Put simply, a blockchain consists of a linked chain that stores auditable data in units called blocks, where each block contains the data, its own hash value (a unique cryptographic value) and a pointer to the hash of the previous block
- Transactions are stored as immutable records in a distributed stateful ledger (timestamped record-keeping in chronological order)
- On a blockchain network, transactions are automatically validated and stored without the need for a centralized administrator
- There are three type of blockchain networks:[5]
  - Permissionless: Permissionless networks are those that anyone can join at any time, such as Bitcoin or Ethereum
  - Permissioned private: permissioned private networks consist of a consortium of finite and well-defined entities that deploy, run and maintain all of the nodes. Generally, these networks are developed, and even maintained, by a blockchain service provider. Examples: the IBM Food Trust
  - Permissioned public: with permissioned public network, entities initiate a network and allow certain parties to join, if they meet certain requirements, such as being authenticated and compliant with regulations. In these networks, the consortium is self-sufficient and does not need to rely on a vendor. Examples: Alastria in Spain, led by an association of over 500 members; the European Blockchain Services Infrastructure (EBSI) in Europe led by the European Union; and LACChain in Latin America and the Caribbean, led by the Inter-American Development Bank and its partners in the programme; and the blockchain network of the Energy Web Chain by the Energy Web Foundation (EWF) consortium

**Smart contracts**

- Self-executing agreements that are triggered based on predefined and agreed events without human intervention
- Since smart contracts run on the blockchain, they allow tamperproof (barring 51% attack) and automated execution of the parameters once data is received
- Represent deterministic automation of a business process involving multiple parties/ users; neither party can default on its obligations and get away with it or tamper with the process once the smart contract is executed

**Oracles**

- Oracles link between the physical world and the blockchain – middleware infrastructure that connects blockchains and any external off-chain system – responsible for reading/ writing data from/to another legacy system and blockchain
- Any number of oracles (nodes) can be combined to make up a decentralized oracle network, which can aggregate the responses of each node in any manner (mean, median, mode, weighted etc.) to ensure there is no single point of failure in the delivery of data, improving data integrity – and thus providing liveness guarantees (that every input is guaranteed to generate an output) and data manipulation protection
- May also have more advanced functions such as providing cryptographic proofs (verifiability of the data), hardware modifications (using trusted execution environments (TEE) to compute transactions), computation (producing privacy, scalability) etc.

**Distributed ledger technology (DLT)**

- An overarching term that encompasses the family of blockchain technologies, including both permissionless and permissioned blockchains, smart contracts and oracles. It also includes non-blockchain architectures such as directed acyclic graphs

**Abstraction layer**

- Abstraction, as a computer science concept, is an approach to preserving information that is relevant in a given context and forgetting information that is irrelevant in that context. It is a process of ignoring temporal details or attributes in the study of systems to focus attention on details of greater importance. For the limited context in this paper, an oracle network that sits between all blockchains and legacy systems to pass data between the two seamlessly and securely is being referred to as an abstraction layer

3

# Interoperability – fundamental pillars

Approaches to develop legacy-DLT interoperability should be open, universal and emerge bottom-up from users rather than being imposed as a top-down standard.

Interoperability in the DLT space is a broad concept. Much research has been undertaken towards enabling blockchain-to-blockchain interoperability, which is implemented largely through inter-ledger transactions. These transactions are intended to allow for digital asset exchange between DLT platforms, with no data or assets from non-DLT systems entering into the overall system.[6] This paper, in contrast, focuses on the exchange of data and dependent transactions between DLT systems and legacy systems.

The primary and predominant reason for DLT-legacy interoperability is to enable smart contracts (on-chain) to use an oracle to fetch information from a legacy system (off-chain), format it, validate it and store it on the blockchain where it can be used to trigger some type of agreement. The reverse use case also exists, whereby on-chain information or some type of command from a smart contract is sent to an external system that uses it for further processing or to act in the real world.

Ultimately, the full power of legacy-DLT interoperability is unlocked when abstraction is used to blur the boundaries between the DLT and the legacy environments.[7] In the legacy-DLT context, abstraction goes beyond blockchain middleware for data retrieval and exchange. It also includes security properties. This level of confidence that cross-network (i.e. across legacy and DLT networks) processes will execute as written with a high degree of security and reliability will allow new use cases to be experimented with. These solutions are referred to here, in this paper, as middleware (explained in more detail in the next section).

FIGURE 5 **Connecting smart contracts with external legacy systems through interoperability standards and legal frameworks**

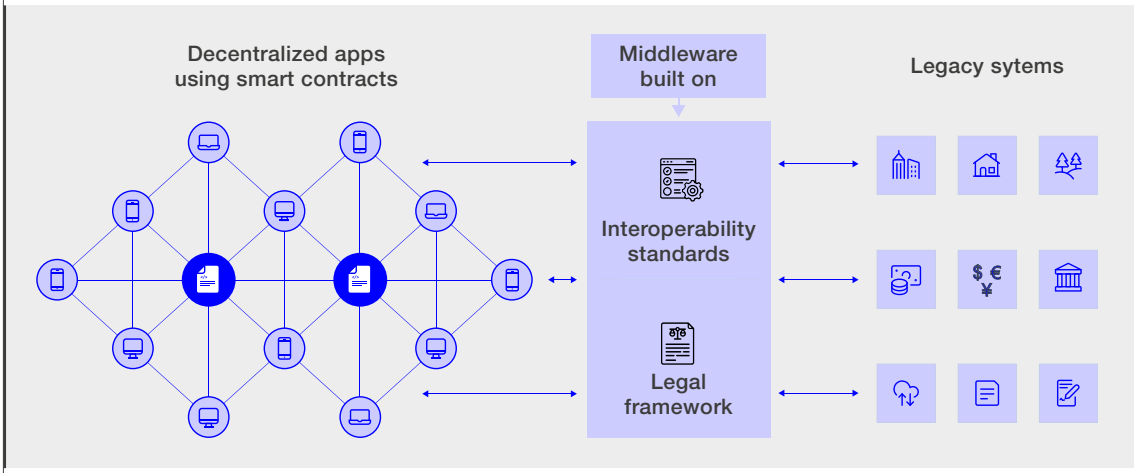


As represented in the figure above, the basic rules for interoperability and legal principles should be established if middleware solutions are to be developed that will allow data/transaction exchange between legacy and DLT solutions in compliance with existing jurisdictional laws. And as these solutions can be developed by different players on different technologies, the underlying interoperability standards should be flexible and comprehensive to work across legacy systems and DLT solutions, and should ideally have the following characteristics:

- Supported by a large open-source community building on a secure middleware
- Support most blockchain networks and legacy systems globally
- Embed protocols and technologies to verify data integrity and counterparty performance
- Use a generalized middleware that can satisfy a wide variety of security and performance needs across all IT systems

Further, it is important to mention that an approach for interoperability should emerge from the relevant stakeholders and users after experimenting with various models. Imposing an interoperability standard from the top down may create negative effects on the blockchain trilemma as it may reduce decentralization by accepting (centralized) data coming from outside. This would also create security loopholes. When standards are used to force interoperability, they may lock all market players into an inferior technology, as Jean Tirole underlined in a ground-breaking article.[8] Accordingly, this paper outlines approaches (strategic pathways) that organizations can adopt to arrive at an ecosystem-driven interoperability solution rather than prescribing a definitive set of standards.

4

# The importance of an open-source, decentralized data validation middleware for enterprise systems

Enterprises running large legacy systems are the key to accelerating the adoption of DLT-based smart contracts.

The fast pace of ongoing technology developments in the DLT space may overwhelm any enterprise or government agency. At times, it may seem like the existing system is already becoming outdated. However, doing a fundamental revamp or even replacing an existing system with a new system may be an expensive and impractical strategy. As a general matter, any organization that is not developing a roadmap on how to integrate future DLT capabilities risks falling behind on innovation that could potentially be relevant for its industry.

A sensible approach that helps to mitigate technology risk and switching costs is to adopt an open-source secure blockchain middleware, wherever the use of blockchain is relevant, that provides legacy systems with universal access to current and future blockchain environments without needing to restructure or rebuild any mission-critical internal backend infrastructure. In addition to universal connection, blockchain middleware might provide some legacy systems with an "in-house" method of developing higher security and reliability standards onto their existing systems using various validation and filtering techniques such as decentralization (redundant confirmation of the same data point in a trust-less manner), reputation systems (filter oracles based on performance quality) and technologically enforced financial incentives/penalties (stake capital to back the quality of oracle services). However, such decentralized oracles may run the risk of majority attacks (where more than 51% of oracle nodes collude to provide incorrect data), single source failure (in the case of a single source of data) and failure of on-chain data confidentiality where additional technological considerations are not taken into account. The sections below discuss some of the mechanisms to manage these risks.

Enterprise systems can support the development of new or innovative services in an ecosystem if those systems can interact with other platforms and networks. This stands true for legacy-to-DLT interoperability as well. However, enterprises and governments may innovate with DLT-related services at a slower pace if they are unable to connect legacy and DLT systems securely and effectively. The ability to easily integrate with any existing and/or future blockchain network and maintain strong guarantees of integrity and determinism will provide governments and enterprises with the ability to build across and interact with counterparties operating in any blockchain environment. Furthermore, organizations or regions may use different blockchain networks, and integration and interoperability among them will be valuable.

This reality presents an immediate need for open-source blockchain middleware solutions that are able to connect information and transactions across different blockchain and legacy systems. By serving as a universal communication layer, this middleware also acts as a common software repository where blockchain developers can share resources such as documentation, technical walkthroughs and information explaining how other users can interact with their system. A standard, open medium for blockchains and systems to connect with one another without permissions avoids integration bottlenecks such as requiring permission from administrators or needing two development teams to develop specific documentation for each new use case.

In the coming sections, this paper attempts to provide strategic pathways towards these goals of building interoperability between legacy and DLT systems. Summarized, the proposed pathways are as follows:

FIGURE 6 **Summary highlight of the interoperability pathways recommended in the paper**

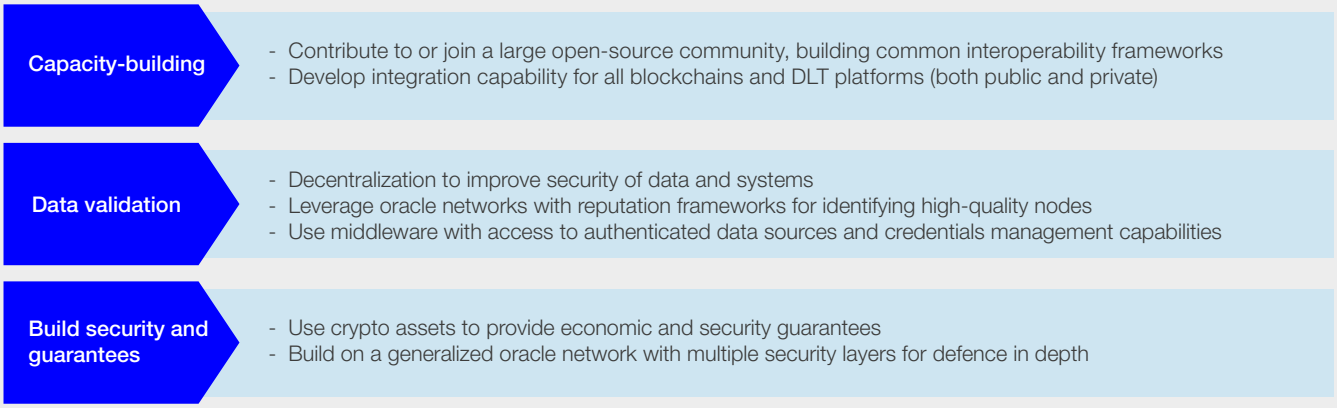

## 4.1 Contribute to or join a large open-source community building on a common interoperability framework

Much like the internet, collectively adopting a standard open network benefits everyone by creating much larger network effects, which in turn creates more value for each individual user. Doing so requires a permissionless environment that anyone can build on, access and use in whatever form they need for their specific use case. Interoperability cannot reach its full potential without being an open network, as it will hit a development choke point if all connection points must rely on a single permissioned middleware. A closed network may also be subject to political stand-offs whereby certain enterprises will not join consortium networks operated by competitors or will refuse to write documentation for certain competitors, thereby limiting access or causing delays. Such an approach will likely result in a fragmented, unscalable system of many different intranets, instead of the one global internet we all enjoy today.

A standard open-source framework provides several advantages. First, blockchain developers need to provide only one set of documentation describing how other systems can talk to their blockchain using secure middleware. This allows any external system to operate on that blockchain (via the middleware) in a permissionless manner by simply following the documentation. In contrast, closed-source middleware is extremely difficult to scale and requires an extensive amount of extra time/resources as every single integration requires communicating with a select few administrators to first get permission to integrate and then to write/edit specific documentation.

Second, open-source networks are far better suited to interoperating among multiple different stakeholders as no one party gains an unfair advantage from owning the IP or sole development rights. Users feel more assured by being able to see and verify the codebase they are using to secure large amounts of value, knowing their competitor doesn't have a special backdoor/privileged access into it. This approach also allows for easy modifications and improved security, as open-source code means more developers can work on the same codebase, pruning any bugs and expanding capabilities that benefit everyone equally.

However, such an open approach doesn't have to come at the expense of regulatory frameworks, as governments can build those on top as different pluggable solutions to ensure compliance based on their own local laws. In much the same way thatas the internet is a single network that is modular and can be adapted to fit certain legal requirements, an open interoperability network can simultaneously support multiple legal frameworks without cross-dependencies on other localities.

**Strategies to adopt open-source culture and community**

Cultivate solutions based on an open-source technology that is generalized and flexible enough to accommodate current and future needs and is permissionless in terms of integration. This means it must be upgradeable and not enforce a particular development framework or technology solution on users.

Use tools and platforms to build open-source communities of developers, researchers, enterprises, users, governments, node operators and other relevant actors to participate in open discussions and express their unique needs. Write documentation on how to interact with the middleware, which then allows everyone to connect to their systems easily and uniformly.

Facilitate government adoption of time-tested open-source software that has been written and improved upon by a global user base of researchers and developers. Mission-critical software systems become more functional when external applications can build additional features for them and more secure when using a global pool of talent to contribute and review the codebase.

**Case study: Governments joining the open-source movement**

Governments are increasingly using open-source development practices to write their software. For example, the number of government employees using Github (a collaborative platform to write software projects) has continued to increase. Linux, an open-source operating system, powers many government servers around the world, including the US Army, which has the largest installed base for Red Hat Linux. Recently, the Government of India put out the codebase of its contact tracing app Aarogya Setu, which has more than 130 million users, in open source.

FIGURE 7 | **Government organizations on GitHub**

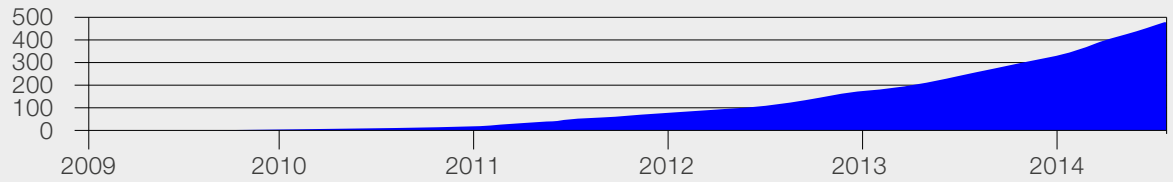


## 4.2 Adopt an oracle network that operates across all blockchain environments

To fast-track the ability for: (1) data providers to share data and services across all of the various blockchain environments; and (2) legacy systems to operate on any chain, limiting the upfront set-up requirements and costs is vital. If data providers and legacy systems have to spend time and allocate resources to integrate separately with each new blockchain environment, they are likely to be very slow to bring their data and services into blockchain ecosystems.

Multiparty business processes may be difficult to transfer on to DLT platforms unless the various distributed counterparties can all support multiple different blockchain environments, allowing for accommodation to partners' preferred platforms. If certain blockchains are not available, there will be gaps in operations that affect the ability of other systems to properly interact: e.g. the supply-chain solution can't easily talk to the finance system, which then interferes with the trade finance network.

Having blockchain middleware running across all blockchain networks allows enterprises to synchronize blockchain applications, legacy tools and internal databases into one trackable and reliable operation. It also provides government entities running legacy digital infrastructure with the ability to rapidly scale up adoption of public programmes and deploy compliance standards across different blockchains through a single gateway.  andofingingThe collective use of a common middleware by users from different blockchains, enterprises and governments reduces the costs of providing and accessing data and services for everyone – achieved via shared financial support of the node operators running interoperability infrastructure. This benefits blockchains through the cultivation of more data-rich environments for developing applications that interface with legacy infrastructure, while equally benefiting legacy systems that can now provide data and services to users across all kinds of DLT networks.

| Strategies to integrate multiple DLT/network capability | | |
|---|---|---|
| Have an oracle network that can run natively on each blockchain so that oracle services are subject only to the throughput and security of a blockchain. This allows blockchains to customize their oracle solution to fit that blockchain. | Implement open-source blockchain middleware that sits as an abstraction layer above all of the blockchains, providing a universal gateway for data providers/node operators to transact with all chains from a single framework. This reduces friction for data providers to set up on different chains, lowering the costs for everyone. | Use that same abstraction layer to allow enterprises to connect to all of the different chains and other legacy systems from a single integration. This vastly reduces the integration work with external systems and consolidates internal operations by ensuring that all of the companies' systems are synced up together. |

FIGURE 8 | **Case study: Interoperability among multiple blockchains – Cosmos and Polkadot**

**Case study: Interoperability among multiple blockchains – Cosmos and Polkadot**

Polkadot and Cosmos are two projects working on interoperability between blockchains.

Polkadot connects several chains together in a single network, allowing them to process transactions in parallel and exchange data between chains with security guarantees. It uses a network of heterogeneous blockchain shards called parachains. These chains connect to and are secured by the Polkadot Relay Chain. They can also connect with external networks via bridges.

(Source: https://polkadot.network/Polkadot-lightpaper.pdf)

Cosmos's architecture is based on a "hub-and-spoke" system whereby a series of different blockchain chains connect to a "central" hub by means of inter-blockchain communication. The goal of Cosmos is to break the barriers between blockchains by allowing them to transact with each other. This vision is achieved through a set of open-source tools such as Tendermint (a byzantine fault tolerant [BFT] consensus algorithm), the Cosmos SDK and IBC, which allows building interoperable blockchain applications. Note that Cosmos is an open-source community project.

(Source: https://cosmos.network/intro)

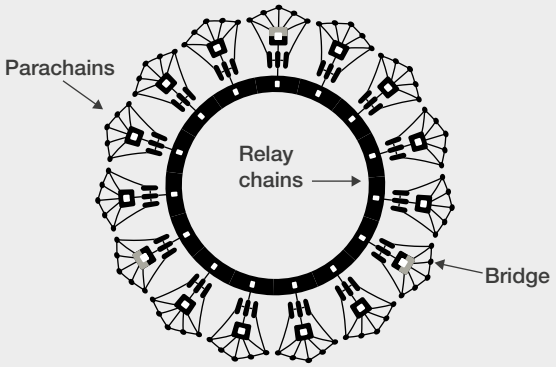


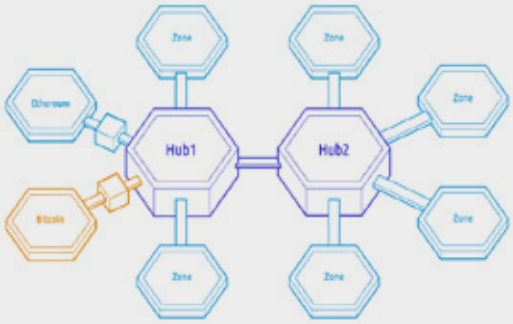

# 4.3 Use decentralization to validate the integrity of data inputs and provide highly available, tamperproof data exchange

Decentralization is one of the core tenets of security for blockchain applications, wherein each node in the decentralized network is financially incentivized to independently validate each transaction based on a common set of protocol rules. In a sufficiently decentralized network with proper incentives for honest reporting, the network's consensus will overpower any transactions that do not comply with the defined rules. This provides network users with deterministic computation (computation which, given an input, will always produce the same output), data integrity and liveness guarantees for the transactions they send to the network for processing.

FIGURE 9 **Two-level distribution of data sources and nodes for oracles**

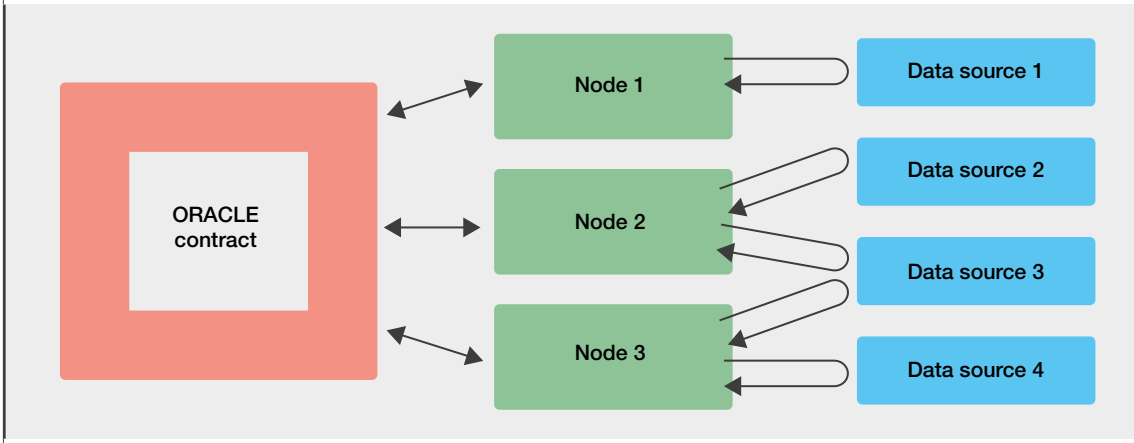


source: https://link.smartcontract.com/whitepaper

Decentralization of oracles should occur at both the node and data source levels (as shown in Figure 2), without compromising the quality of any one component (e.g. incorporating low-quality data), to ensure there are multiple layers of redundancies. Decentralized oracle networks remove any single point of failure in the delivery of data to prevent the smart contract from relying on any one single node or source of truth.

For further guarantees about the integrity of data, especially in situations of a single data source, specialized cryptographic proofs can be provided that verify the authenticity of the data point's origin. This reassures users that external data came directly from a particular web server and was not tampered with en route, as only the specific web server can provide a unique signature proving its origin. By combining this cryptographic technique with decentralization (e.g. multiple nodes and multiple data sources), the inputs and outputs of a digital agreement can become as tamper-resistant as the smart contract itself.

For governments, this becomes especially imperative where citizen services and claims can be verified through an independent validation system that records their request on a DLT, and the governments/enterprises servicing that request are bound by non-repudiation (assurance that entities cannot deny any information they have signed and provided). When combining decentralization with high-quality nodes filtered through a reputation system (described below), citizens can receive stronger guarantees on certain government services by having them validated by a decentralized network, while governments can reduce manual and complex coordination processes in relation to information and asset transfers by offloading them to a reliable decentralized network of oracles.

## Strategies to develop decentralized data validation

Apply the security model of decentralization to oracles relaying data between on-chain and off-chain systems. Decentralized oracle networks use multiple independent/Sybil- resistant nodes (mechanisms preventing nodes creating multiple identities from influencing the network) to validate the same off-chain data point, providing liveness and preventing a single oracle from corrupting the data.

Aggregate data from multiple high-quality data sources, retrieving the same data multiple times asynchronously and/or incorporating dispute resolution processes to safeguard against a single data source or single API call being the sole arbiter of truth for a smart contract. This approach provides multiple layers of redundancy and ensures that if a data provider or node goes offline or fails to deliver the data, the overall data feed is still reliable.

Select a multitude of independent high-quality oracle node operators that have been security reviewed and meet the collective standards set out by the stakeholders involved, such as a strong performance history of delivering reliable services, known registered entities run by experts such as leading DevOps, and any other consideration deemed important.

### Case study: Crypto price referencing using decentralized oracle networks: Chainlink and MakerDAO

Chainlink, a framework for decentralized oracles, supports decentralized oracle networks that feed tamper-proof price reference data to smart contracts on-chain. Similar to blockchains, which use decentralized computation to create data integrity, oracle networks such as Chainlink employ the same decentralization model to aggregate price data and store it on-chain, ensuring it is highly accurate, available and resistant to manipulation. Using a decentralized network of security-reviewed oracle nodes and data sourced from off-chain data providers, Chainlink maintains an on-chain price feed that updates every time there is a small deviation from the latest price stored on-chain. Blockchain applications can read the on-chain price data and use it to execute automated functions, such as ensuring a loan is fully collateralized or settling a derivatives contract.

(Source: https://chain.link/)

MakerDAO, a decentralized stablecoin protocol, uses oracles to fetch asset prices of supported collaterals. The reference price is calculated via a medianizer – a smart contract that collates price data from several external independent price-feed operators. These operators monitor the asset prices from external sources. Third-party relayers are then used to forward their prices on-chain.

(Source: https://makerdao.com/)

FIGURE 10 | **A visualization of the decentralized oracle network supporting the Chainlink BTC/$ price feed**

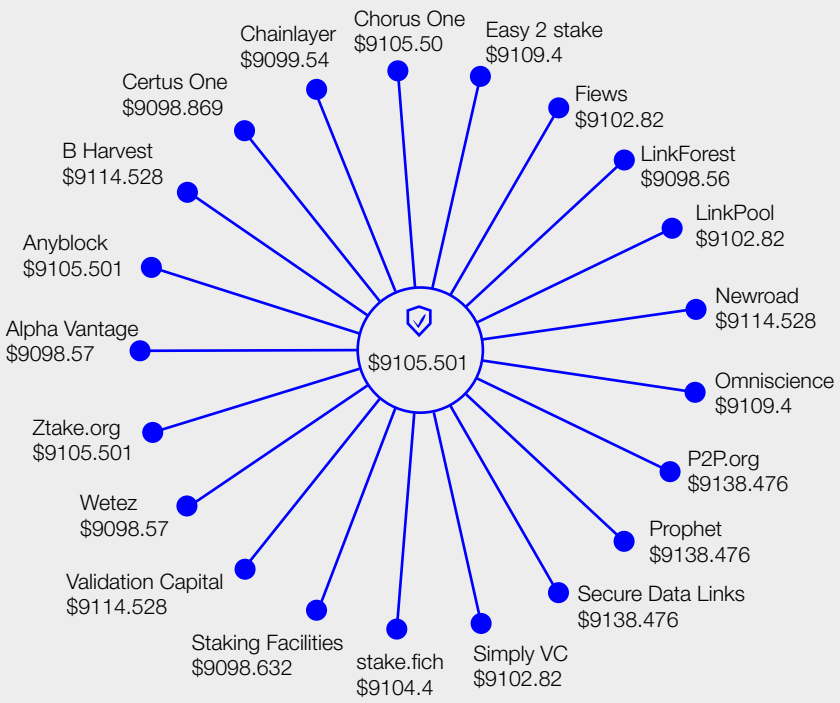

# 4.4 Demand oracles that support access to authenticated data sources and credentials management capabilities

Smart contracts use data to directly execute automated business processes without a human intermediary. Thus, data quality is an extremely critical component in determining the smart contract's security and reliability to avoid a "garbage in, garbage out" scenario.

Data quality may take capital to create – often the result of large network effects, unique business processes and/or an intelligent team of data analysts. Incentives must exist to maintain the production of high-quality data, meaning that there must be a framework for handling credentials (login, passwords) to limit access to paying users for high-quality authenticated data APIs. Without having credential management capabilities built into an oracle system, data is at risk of being pirated, thus limiting the data provider's ability to generate revenue and continue generating high-quality data. Generally, free data APIs lack incentives to keep data up to date and machine-readable, making them potentially unreliable and insecure as inputs to trigger the movement of large amounts of value.

Data providers have two primary options for how to make data available to blockchain networks: provide data to an existing oracle network and/or run their own oracle node to provide origin-signed data directly to smart contracts. Through an existing secure and reliable oracle network, data and API providers do not have to change anything about their existing business model. The oracle nodes can pay the data providers in fiat currency using typical payment plans in use today.

Alternatively, data and API providers can run their own oracles to both sign data with their private key and sell directly to blockchain-based smart contracts. This provides them with a direct method of participating in the blockchain ecosystem by earning additional revenue from selling directly into new blockchain markets, as well as adding additional security to their data by cryptographically signing it when broadcasted to a blockchain. Governments and enterprises can benefit from such a method by running their own oracle nodes, providing users with strong guarantees that the data and services they provide came directly from official channels.

| Strategies to integrate high-quality data sources | | |
|---|---|---|
| Develop/integrate highly secure credential management capabilities to access high-quality data providers and enterprise systems such as web APIs, internet of things (IoT) networks, CRM/ERP systems and various other legacy systems that require authorized logins. | Use decentralized oracle networks to improve data quality by aggregating data from multiple data sources to provide a single source of data that is likely to be accurate. Aggregation techniques can be customizable (average, weighted etc.) and performed cost-effectively off-chain with supporting cryptographic proofs that are provable on-chain. | Use oracle networks that keep on-chain data fresh and reflective of real-time conditions in an automated manner while retaining key security properties. Oracle networks should also be flexible to support on-demand requests for various types of off-chain data and API services. |

**Case study: Use of APIs**

Making real-world data available on-chain requires oracle networks that are compatible with the existing API infrastructure already widely used throughout global economic systems by businesses, enterprises and governments. To ensure an accelerated and seamless transition, this means oracle networks need to support access to credentialed systems without any changes to the current API systems, as well as a framework for API providers to run trusted oracles in order to sign their own data and sell it directly. A robust blockchain oracle solution should enable data providers to both sell their APIs to oracle networks and publish signed data directly on-chain without any changes to their existing systems.

FIGURE 11 | **A visualization of predominant API ecosystem**

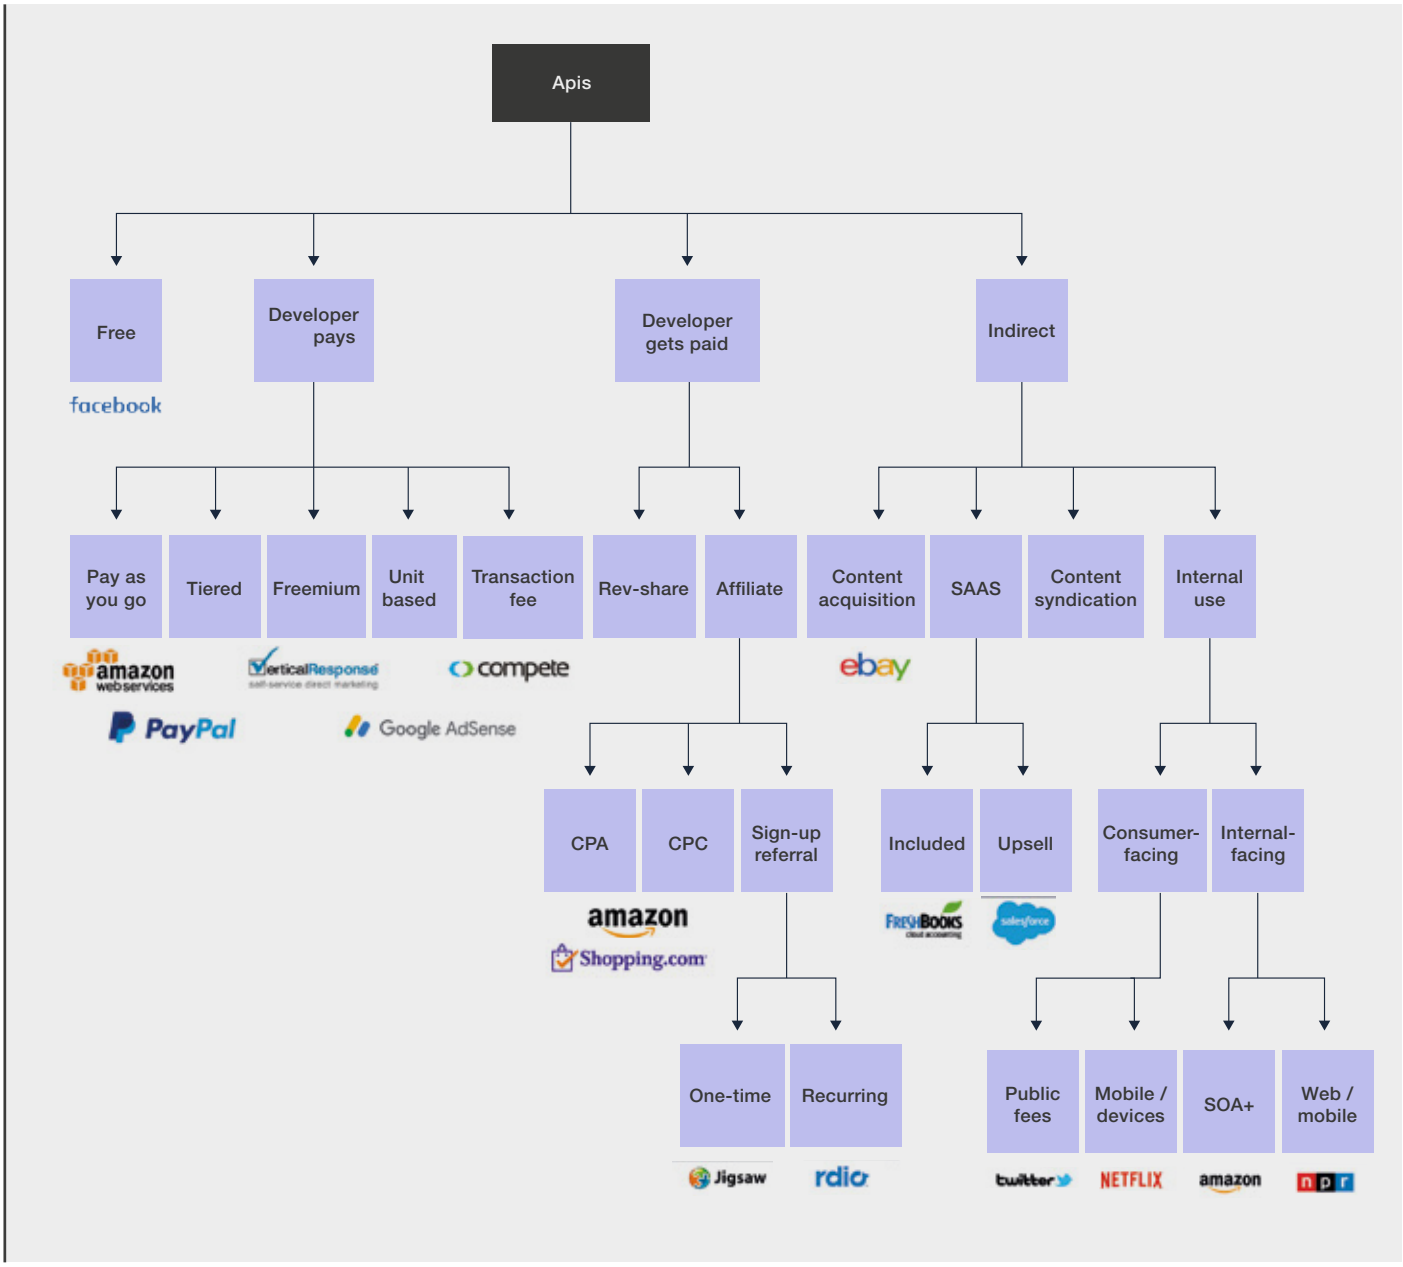


**Case study: Ramp Network uses open banking APIs to execute crypto-to-fiat transactions**

Ramp Network is one of the licensed players under the second Payment Services Directive (PSD2) for open banking (European banks). In order to provide peer-to-peer, crypto-to-fiat transactions (using atomic swaps), it uses a payment oracle to verify a buyer's payment and send a proof-of-payment to the network. Here, the data source is single (just the user's bank) and privacy requirements are higher to ensure the user's identity and transaction history aren't made public.

(source: https://swaps.ramp.network/)

## 4.5 Use crypto assets to provide for crypto-economic guarantees

If smart contracts are to automate business processes between multiple parties, then there need to be very clear terms between the smart contract and the oracle mechanism responsible for connecting it to all the inputs and outputs in the form of a binding service level agreement (SLA). These terms need to be supported by economic incentives and penalties, both of which are enforceable on the blockchain. By tying the security and reliability of oracle infrastructure directly to technologically enforced financial incentives, oracle providers are forced to have "skin in the game" that is both rewarded or punished financially based on the quality of their performance (as per their SLAs). The collateral staked by oracle nodes acts as a crypto-economic guarantee of optimal performance as nodes have a direct financial stake in the correct functioning of their oracle services. The oracle's performance data needs to be recorded in a tamper-proof manner and available to other potential data requesters (as part of the public reputation system mentioned below).

Off-chain legacy systems, data providers and oracle node operators need to be held to the same standards as blockchain systems to ensure they are deterministic and reliable enough to be trustworthy to deliver the data that will ultimately be used to execute the smart contract. Through the SLA framework, legacy systems become trusted oracles to smart contracts, as these agreements are binding and technologically enforced with financial and reputational rewards/penalties based on the quality of performance. Moreover, the outcomes of SLAs can be stored as immutable historical records of a node's performance within reputation systems (described in the next section).

Open-source public, permissionless blockchain networks provide crypto-economic security using a native token, which is used to stake as collateral and pay for network services. Bitcoin and Ethereum are both examples of open-source, permissionless blockchains that have found success through having a native token in order to fund the miners who secure the network. Thus, the token's value is tied to the overall security and reliability of the network, creating a positive feedback loop of incentives. By doing so, blockchains become public goods that are collectively secured and maintained by a public community of users and various stakeholders rather than a centrally operated, for-profit entity. Open-source oracle networks follow the same public good model, where secure oracle services are fuelled by a native token adopted by both oracle node operators and end users.

For organizations and governments using permissioned blockchain networks, users can employ third-party services that allow them to pay in whatever currency they want, yet still ensure oracles are paid in the native token on the backend. This can be done in a manner similar to how invoicing works today, where a third-party relayer collects user fees in any currency, converts those currencies to native tokens, and pays the node operators in the native token, taking only a small commission fee for their services. In this regard, the token is abstracted away from governments and enterprises ever having to use it, yet such institutions can still participate in public networks with a native token and benefit from crypto-economic security.

| Strategies to provide economic and security guarantees | | |
|---|---|---|
| Create a service level agreement (SLA) between the smart contract requesting off-chain services and the oracle providing those services. The service agreement needs to be enforceable directly on-chain based on very clear pre-agreed upon and digitally signed terms between the two parties. | Adopt a framework where oracles deposit collateral as an economic guarantee to back their services, which is either returned (with a commission for performing work) or taken as a penalty for not performing according to the terms of the service agreement (the oracle node went offline or the node provided data that deviated too far from the other nodes' responses in a decentralized oracle network). | Record the performance data of the oracle node on the blockchain and feed it directly into reputation systems. This allows future customers to determine the quality of the oracle node operators and enables existing smart contracts to potentially remove nodes from data requests that were recently reported to have been malicious (as per the terms in the SLA) or unreliable. |

**Case study: How does the Telecom Regulatory Authority of India (TRAI) implement financial disincentives with a DLT-based audit trail and without the use of crypto assets??**

The Telecom Commercial Communications Customer Preference Regulations, 2018, by TRAI in India, uses a DLT-based audit trail and consent to formulate financial disincentives to the originating access provider (OAP) in order to curb unsolicited commercial calls sent through its network. Excerpt: "I"f OAP fails to curb UCC, Financial Disincentives for not controlling the Unsolicited Commercial Communications (UCC) from RTMs by the access provider in each License Service Area for one calendar month shall be as under: -

| | Value of "Counts of UCC for RTMs for one calendar month" | Amount of financial disincentives in Rupees |
|---|---|---|
| (a) | More than zero but not exceeding hundred | Rupees one thousand per count |
| (b) | More than hundred but not exceeding one thousand | Maximum financial disincentives at (a) plus Rupees five thousand per count exceeding hundred |
| (c) | More than one thousand | Maximum financial disincentives at (b) plus Rupees ten thousand per count exceeding one thousand" |

Here is an interpretation of the penalties rule: In one year, telemarketers are allowed a total of 12 violations with an accompanying penalty. After 12 violations, telemarketers will be blacklisted for two years. Violations are tracked via the DLT through a complaint and entity module, which can be referred to by TRAI to deduct the appropriate penalty from the deposit amount and provide notice to the telemarketer. Telemarketer and telecommunication service providers (TSP) must maintain a registration deposit with TRAI with a minimum balance. As deductions are made from it because of violations, telemarketers or TSPs have to top it up with additional deposits. Failure to do so will lead to the issuance of a notice to pay or, if payment is not forthcoming, placement on a deny list.

## 4.6 Leverage oracle networks with marketplaces and reputation frameworks to identify high-quality node operators

Just like consumers want to research the quality of a doctor or ask their peers about the reputation of a business, users need to be able to determine the quality of oracle node operators, whether that be independent nodes or legacy systems and data providers operating as trusted nodes. It is risky and unlikely for a smart-contract creator to entrust the security of a large amount of value to an anonymous or unproven node operator with no provable or verifiable metrics about its ability to provide high-quality services and data integrity.

Independent reputation systems for oracles can take many forms. One form is simply having third-party services provide certifications of node operators by undertaking security reviews of their infrastructure, performing know your customer (KYC) protocols on the owners, issuing regulatory certifications and various other types of approval processes. These certifications can be displayed in online marketplaces where nodes list their services, which not only allows nodes to highlight their merits but also provides users with an interface to filter nodes according to the features they deem most important. This approach can come in handy for users who wish to meet certain General Data Protection Regulation (GDPR) requirements (described further in the section: “Addressing legal and data privacy concerns”).

Another form is using historical transactional data from the open-source blockchain as a means of determining the performance of an oracle node operator. Since nodes sign the data they provide with their unique private key, their performance can be tracked by third-party providers and sorted by various metrics such as number of successful jobs, average response time, quality of clients served and more. The SLA framework can feed right into reputation systems to provide even more advanced metrics, for instance,such as the amount of crypto-economic security provided, number of penalties for downtime, total penalties for bad data outliers and any number of metrics of importance to users. Reputation systems do not have to take any singular form but can involve multiple opt-in reputation systems existing in parallel that vary based on which needs the different stakeholders find most important.

It is very important for governments and enterprise systems receiving data from numerous external systems (nodes) to be able to assess those systems’ integrity and data quality. It’s likely that on-chain performance with certain key certifications will emerge as an important model for choosing nodes. Such a filtering system allows them to identify reliable and compliant oracle node operators; it also encourages external systems to improve their reliability and quality so that they rate higher in the filter system. Through transparency, node operators/ data providers can be held accountable as any poor or malicious performance is filtered out and no longer chosen for future quorums.

**Strategies to assess quality of legacy and operating nodes**

| | | |
|---|---|---|
| Have defined SLAs where oracle nodes sign off on all off-chain data and on-chain transactions they provide as inputs and outputs to smart contracts using their private key. This makes their historical performance as a node operator within well-defined SLAs cryptographically verifiable and immutable on-chain for other users to see. A web of trust framework is created when users can verify the historical performance of a node. | Encourage third-party/open-source reputation platforms that can use on-chain performance data to create evaluation frameworks that rate the quality of nodes based on a number of different metrics – number of successful jobs completed, uptime, the number of applications using a node operator, response times etc. Reputation platforms provide a way of filtering the quality of node operators and create a competitive environment where nodes compete to provide the best-quality services in order to secure more future revenue. | Incentivize listing services for node operators to register their nodes, which provides Sybil resistance, as well as for displaying additional certifications such as third-party security reviews of a node’s set-up, KYC compliance etc. This gives users a reliable place to find node operators that meet their security and service requirements. |

## Case study: Reputation-based technical infrastructure – Amazon Web Services

A reputation-based system is a critical framework for incentivizing performance and reliability across a large, distributed network of operators. Amazon Web Services (AWS) is a cloud-computing provider that publishes up-to-the-minute service availability, as well as performance history via a service health dashboard. AWS's status history log allows users to analyse individual API gateway performance across different regions and periods of time.

(Source: https://status.aws.amazon.com/)

FIGURE 12 | **Case study: Reputation-based technical infrastructure – Status history of AWS infrastructure**

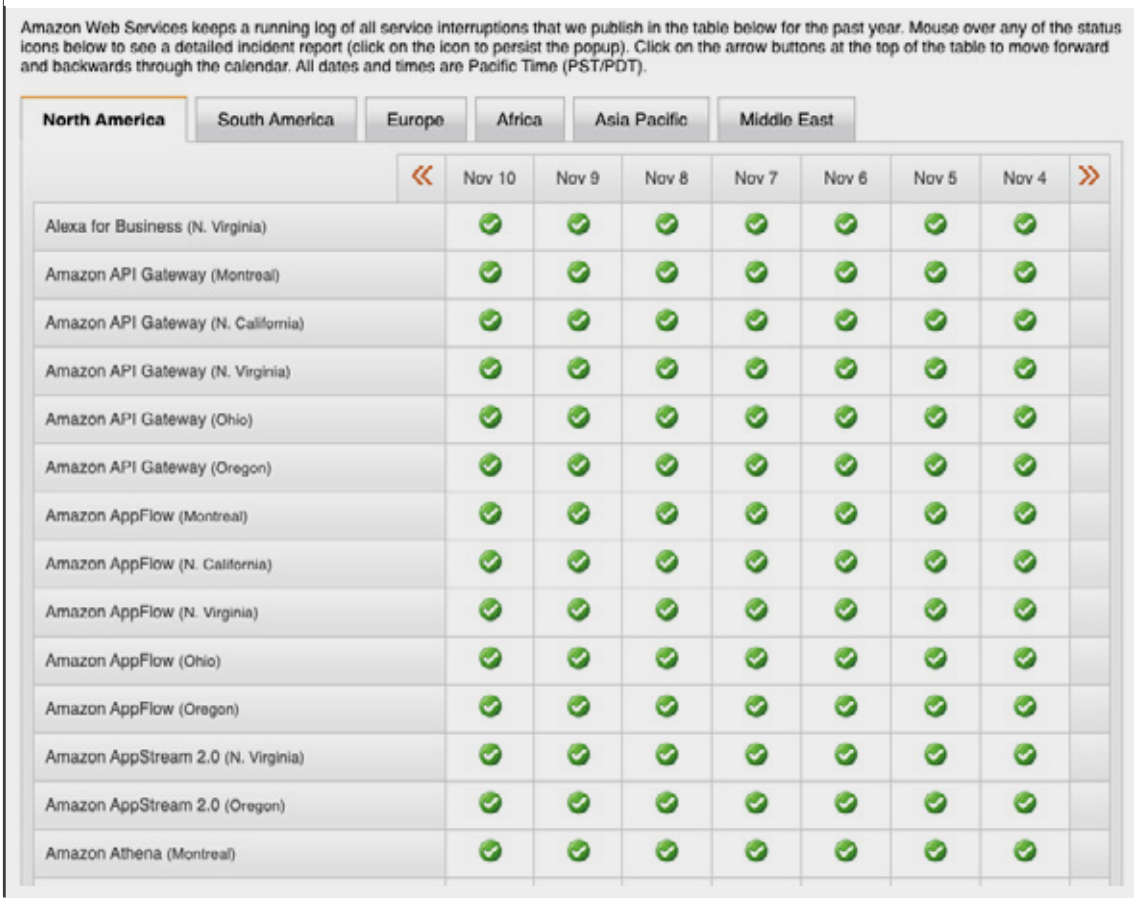

# 4.7 Build on a generalized oracle network with multiple security layers for defence in depth

Legacy-DLT communication channels need to be secure to ensure that any data from the legacy system is not intercepted and corrupted on the way. This is referred to as a man-in-the-middle attack, which becomes even more important in cases where transaction triggers are issued out from decentralized applications to legacy systems (e.g. payment settlement in bank accounts after a transaction happening on-chain).

However, users (governments or enterprises) require different levels of security guarantees for their transactions being fulfilled by smart contracts, which come with different trade-offs and costs. These security expectations and guarantees can vary extensively across industries, use cases and transacting parties. Hence, there needs to be a flexible framework that allows developers to layer on multiple security approaches, which become progressively greater or more specialized depending on the data triggering their smart contract.

Many applications have single-source datasets where data validation through decentralization is not possible. In such cases, users need different forms of security guarantees to prove data and computation integrity when decentralization is not applicable or cannot sufficiently provide enough security/reliability to the contract. Such situations may include using cryptography for single-source data requests, running off-chain computation for business logic concealment, generating on-chain privacy for transaction confidentiality and using trusted hardware for oracle service confidentiality.

## Strategies to execute defence in depth security

Use oracle configurations that provide reliable, cryptographically provable services without exposing the user's data to the node operators. One prime example is a trusted execution environment (TEE) – a trusted computing environment for running code, with a user-triggered attestation that proves the TEE is running as programmed. It provides integrity and confidentiality of the data, the code (running inside the TEE) and the output it generates.

Use oracles that provide cryptographic proofs to authenticate single-source data, such as verifying an SSL/TLS certificate (a transport layer certificate for web transfer of data to ensure identity and protect data privacy) for web data (similar to a website signing data to prove it came from the specified web server).

Use oracle services that create on-chain privacy so that the public cannot see the sensitive data of smart contracts, such as being able to identify payment schedules. This is exceptionally important in the age of AI as people run data analysis on the blockchain to look for patterns and associated addresses.

**Case study: Trusted computation; single-source data authentication; use of TEEs; on-chain privacy of contracts**

**Trusted computation:** On a distributed state machine such as a blockchain, every transaction is executed and validated on every node of the network. It is this redundancy and transparency that provides a network with its integrity but this also comes at the cost of performance and confidentiality. By offloading some work off-chain, participants can trade off resiliency and integrity for performance and confidentiality. The use of "trusted computing" is intended to maintain resiliency and integrity guarantees as much as possible while affording the additional performance and confidentiality. Trusted computation is a framework for optimizing the performance and privacy schema of blockchains while also providing the security and reliability guarantees of a highly decentralized network.

(Source: https://www.hyperledger.org/blog/2019/10/03/introducing-hyperledger-avalon)

**Single-source data authentication:** DECO is a privacy-preserving oracle protocol that allows data transferred over HTTPS/TLS, which is most of the world's data today, to be authenticated as coming from a specific server without leaking any sensitive data on-chain or requiring any server-side modifications. This allows modern internet-connected infrastructure to become validated in a privacy-preserving manner, so that sensitive and confidential data can be attested to, and used by, blockchain-based smart contracts.

(Source: https://arxiv.org/pdf/1909.00938.pdf)

**Use of trusted execution environments (TEEs)**: Town Crier is a TEE-based oracle solution that can verify TLS certificates, proving that the oracle retrieved data from a specific web API.

(Source: https://eprint.iacr.org/2016/168.pdf)

Provable uses TEE environments (such as Amazon's EC2, Google's SafetyNet, Qualcomm's QSEE, Ledger's Nano S and Intel's SGX) to minimize vulnerability. To ensure integrity of the data, Provable uses TLSNotary to digitally sign TLS data from https websites. However, the risk of many data sources being compromised is still a challenge, which is also present in decentralized solutions.

(Source: https://provable.xyz/)

**On-chain privacy:** Mixicles use oracles to split a smart contract up into two parts: the execution of the contract and the resulting output payment. The public is unable to correlate the two parts, resulting in on-chain privacy for financial contracts without changing underlying public blockchain infrastructure.

(Source: https://chain.link/mixicles.pdf)

5

# Addressing legal and data privacy concerns

Organizations intending to make their legacy systems capable of interacting with DLT networks and smart contracts need to be aware of data, privacy concerns and legal compliance.

Enterprises and governments deal in highly regulated environments and need certain frameworks to ensure concerns about data privacy and legal liability are addressed. Thus, having a generalized oracle network that is flexible enough to be able to accommodate many different types of legal requirements is critical to becoming a standard around the world.

## GDPR implications

Since the oracle is handling the data of users, it may come under GDPR guidelines with respect to data handling, data processing and other obligations under the guidelines. Given that blockchain middleware does not enforce specific node selection and allows oracle nodes to run in parallel without cross-dependencies, users can choose the oracle nodes they want to use – which includes filtering them according to geographic region to ensure they comply with GDPR requirements. This allows countries to keep data within their borders while not stifling innovation in terms of using new blockchain applications. If GDPR becomes important for oracle infrastructure, node operators can display certifications to prove they operate within a certain jurisdiction, enablinge enterprises and governments to make informed decisions.

## Data privacy

The other concern often expressed is how to retain data privacy with respect to the oracle itself, which may be handling sensitive data. In this regard, advanced cryptographic techniques can be used that allow the oracle to retrieve data from external sources and feed it to the contract without being able to view the data or know to whom they are sending it. As explained in the "defence in depth" section, some of the specific methods being used include the introduction of trusted hardware in an oracle node set-up to provide confidential computing where not even the node can view the data, or non-hardware techniques such as zero-knowledge proofs, which allow the oracle to attest to the integrity of external data without ever seeing the specific data, exposing the other related data or making it publicly available/viewable on the blockchain.

## Legal liability

Both oracle node operators and data providers can be held financially and legally accountable through binding service agreements, which are smart contracts signed by each party before the data has begun to be handled and delivered. The service agreement contains all of the specific parameters and conditions regarding what happens during edge cases, such as a node going offline during data delivery, manipulated data, detected outliers and more. Remedies can include dropping a node's reputation, slashing its collateral stake, and/or removing it from being selected in future oracle networks. Since these terms are codified as immutable smart contracts on the blockchain with a digital signature for each entity, such agreements can be audited by any party and brought into court if needed during times of dispute, especially when considering users can select known node operators.

It should be noted that smart contracts are not supplanting litigation but are attempting to vastly reduce it through the immediate technological enforcement of terms upon the receipt of data as opposed to probabilistic enforcement by humans with long, drawn-out processes to settle disputes. What is likely to happen is that certain standardized templates will emerge for how contracts are handled and, over time, slight modifications will be added depending on the specific contract. Legal precedents are likely to be set over time when certain situations arise – no different than traditional legal systems today.

6

# ‘Blockchain Abstraction Layer’: A bridge between blockchains and legacy systems built on interoperability fundamentals

Decentralized oracle services can become the abstraction layer for legacy and DLT systems to interact with and unlock hidden value by combining the utility of both worlds.

Similar to how the internet provides universal access to information and systems, smart contracts need a robust, open-source single integration abstraction layer that has all the inputs and outputs readily available to build universally connected contracts. Intranets were not scalable and, similarly, blockchain applications will not reach mass adoption if either systems are not compatible for a smart contract to work or all integration and innovation focuses on a narrow set of networks. By making data, systems and documentation available through a standard permissionless framework collectively maintained by the greater open-source community, smart contracts on any blockchain can be written about any event using any data source and fully integrated with any system, all with substantially less time and fewer resources.

Blockchains, as they are currently architected, are not ready to support all of the needs of legacy systems and are not specifically designed to handle certain processes that run outside of a blockchain network. Therefore, what is needed is a permissionless abstraction layer that anyone can use to connect any on-chain and off-chain systems together and to provide deterministic operations in business processes, without having to completely rebuild the backend of legacy systems. Without deterministically enforced accountability in the performance of off-chain systems, universally connected smart contracts lose reliability once they interact off-chain. An abstraction layer allows legacy systems and blockchains to preserve the properties that make them uniquely valuable (blockchains stay decentralized and completely deterministic while enterprises can continue to benefit from processes that operate better in a centralized and controlled manner). Such a secure and generalizable abstraction layer also acts as a standardized medium for designing how the two environments will interact, in which neither side has any built-in advantages or the ability to tamper with the system for personal gains. This facilitates much-needed trust between the two environments that is equally verifiable by both counterparties.

FIGURE 13 **Bidirectional interoperability bridge schematic between legacy and DLT system**

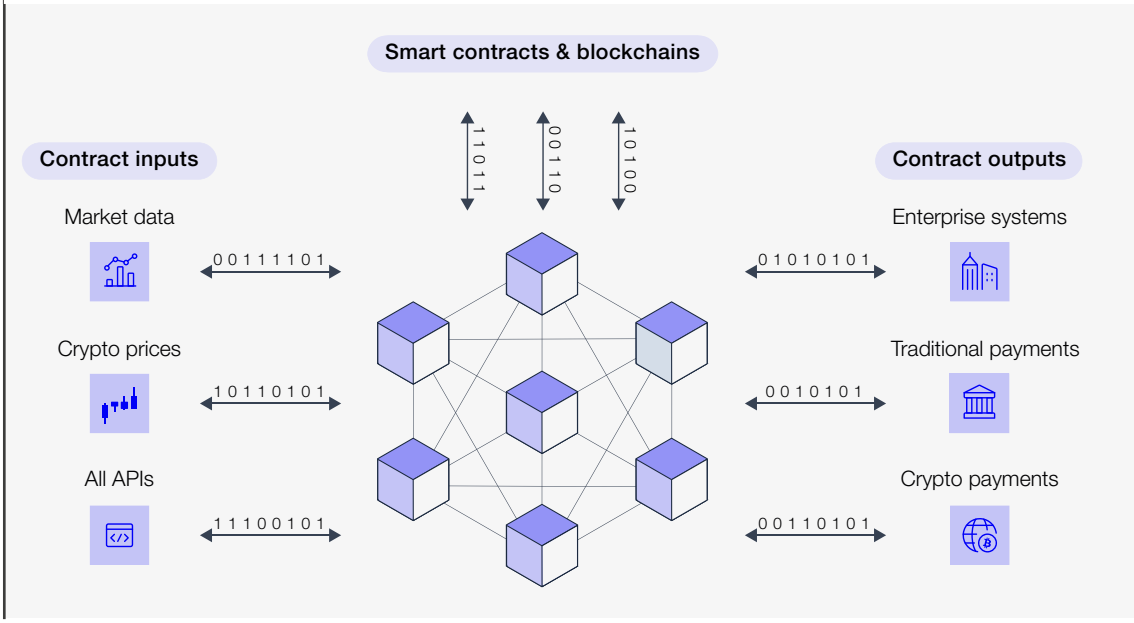

The key features described in the above sections tread the path towards adopting universally connected contracts that use an oracle network which runs natively on any blockchain and is connected to any legacy system. Such a system must have a framework on which new blockchains and legacy systems can be easily onboarded and which is flexible enough to run at the native speeds of different blockchains and off-chain systems. It requires support for bidirectional communication wherein oracles can read data from off-chain networks and write it on-chain (e.g. weather data used to trigger a crop insurance smart contract). Alternatively, oracles must be able to use on-chain data to trigger off-chain actions (e.g. a trade finance contract that settles on SWIFT). The system must use a framework for binding service agreements that outline how off-chain systems will interact with on-chain systems, specifically the terms of the off-chain interaction, such as provide "*x*" data at "*y*" time and if not, then "*z*" will be enforced. It must also provide a framework for crypto-economic security guarantees, such as penalty/deposit systems (staking collateral) and support the use of third-party reputation systems for evaluating different oracle, through immutable on-chain performance data and certification services.

By combining full connectivity between on-chain/ off-chain systems with binding agreements that have security, economid and reputational guarantees, an abstraction layer for building universally connected contracts is created that can automate business processes across any systems in a much more deterministic manner with lower overheads. With such a blockchain abstraction layer, enterprises, governmentd and citizen stakeholders can focus on using that frameworkdand the accompanying open source documentation to create systems of automation that can support any use case or design requirements.

7

# Conclusion: What does an ideal integrated system look like?

Policy-makers and business heads should undertake a projected impact assessment to see how their legacy systems will be transformed by utilizing interoperability pathways prescribed in the white paper.

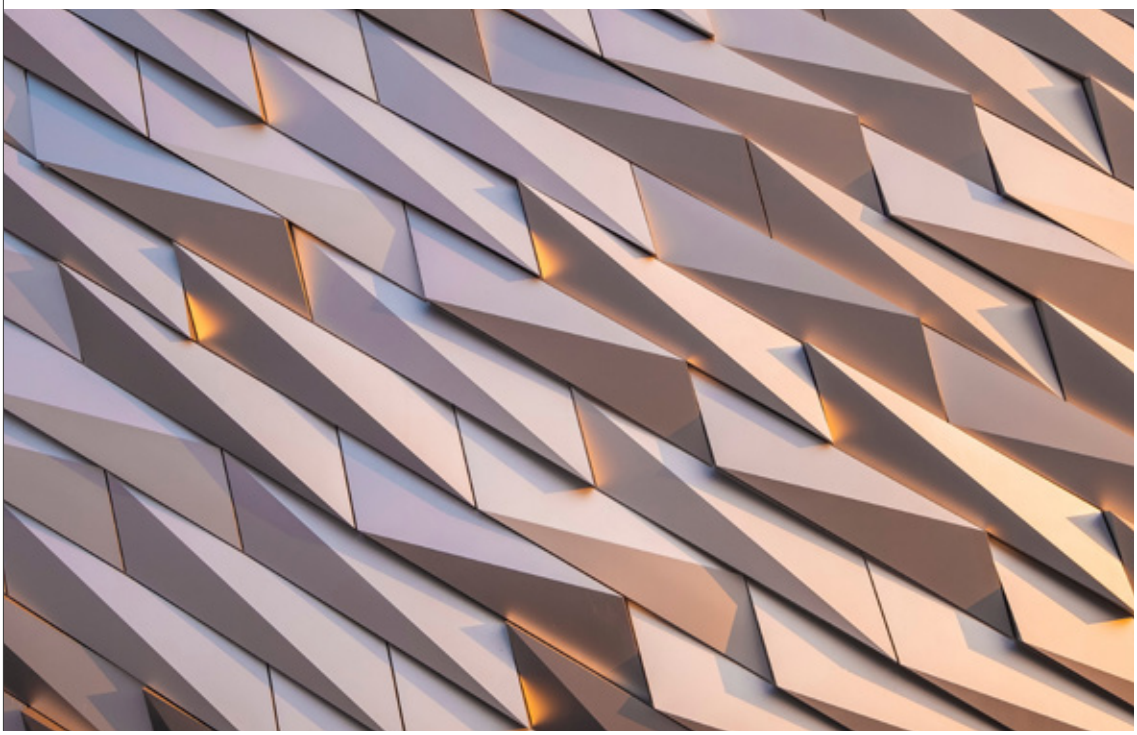

This paper started by discussing the issues that a legacy system faces while dealing with a plethora of external data sources, transaction inefficiencd and multiple human interventions that make the system even more unreliable. For systems that benefit from the use of blockchain technology, the recommendations made in the above sections can be used to build a better, more efficient blockchain-based system with high data integrity. As a final example of the impact of an interoperability bridge, India's crop insurance scheme can highlight the case of performance improvement by adopting the pathways outlined in this white paper.

The national crop insurance portal (NCIP), the Indian government's central legacy application used to operationalize the crop insurance programme, can use all of the interoperability pathways recommended to securely automate business processes with regards to insurance assessments and payouts. Since the insurance company pay-outs are intrinsically dependent on yield assessment of crops through crop-cutting experiments (CCE), which is generally conducted by multiple on-the-ground agencies, validation of CCE data using decentralized oracle networks will ensure the pay-outs are correct as owed to the farmers. The NCIP can use trusted execution environments (TEEs) to make private transaction executions to farmers directly transferring the welfare payments to their bank accounts.

An oracle network with a proven incentive and reputation system would also unlock a marketplace for the larger open source community to enrich NCIP's forecast systems with increasingly localized and granular data. The NCIP must also use high security credentials and APIs to access the high quality and more granular geospatial data from the remote sensing satellites of the Indian Space Research Organisation (ISRO) and Indian Meteorological Department (IMD). Records from various state governments can be securely fed into the portal using adapters and APIs that easily integrate with their existing systems. Furthermore, thie extremely critical weather data can be validated in addecentralized manner (using multiple data sources), and nodes can be adequately incentivised/penalszed based on their accuracy and the integrity of the weather data. Here again, reputation systems and listing marketplaces can help incentivise the generation of better-quality weather information systems. Altogether, this rich ecosystem of real-world inputs supported by an underlying oracle infrastructure provides an innovative yet easy-to-integrate solution for improving the country's crop insurance programme.

FIGURE 14 **How the Pradhan Mantri Fasal Bima Yojana (the Prime Minister's Crop Insurance Scheme) may look like when DLT-legacy interoperability pathways are adopted**

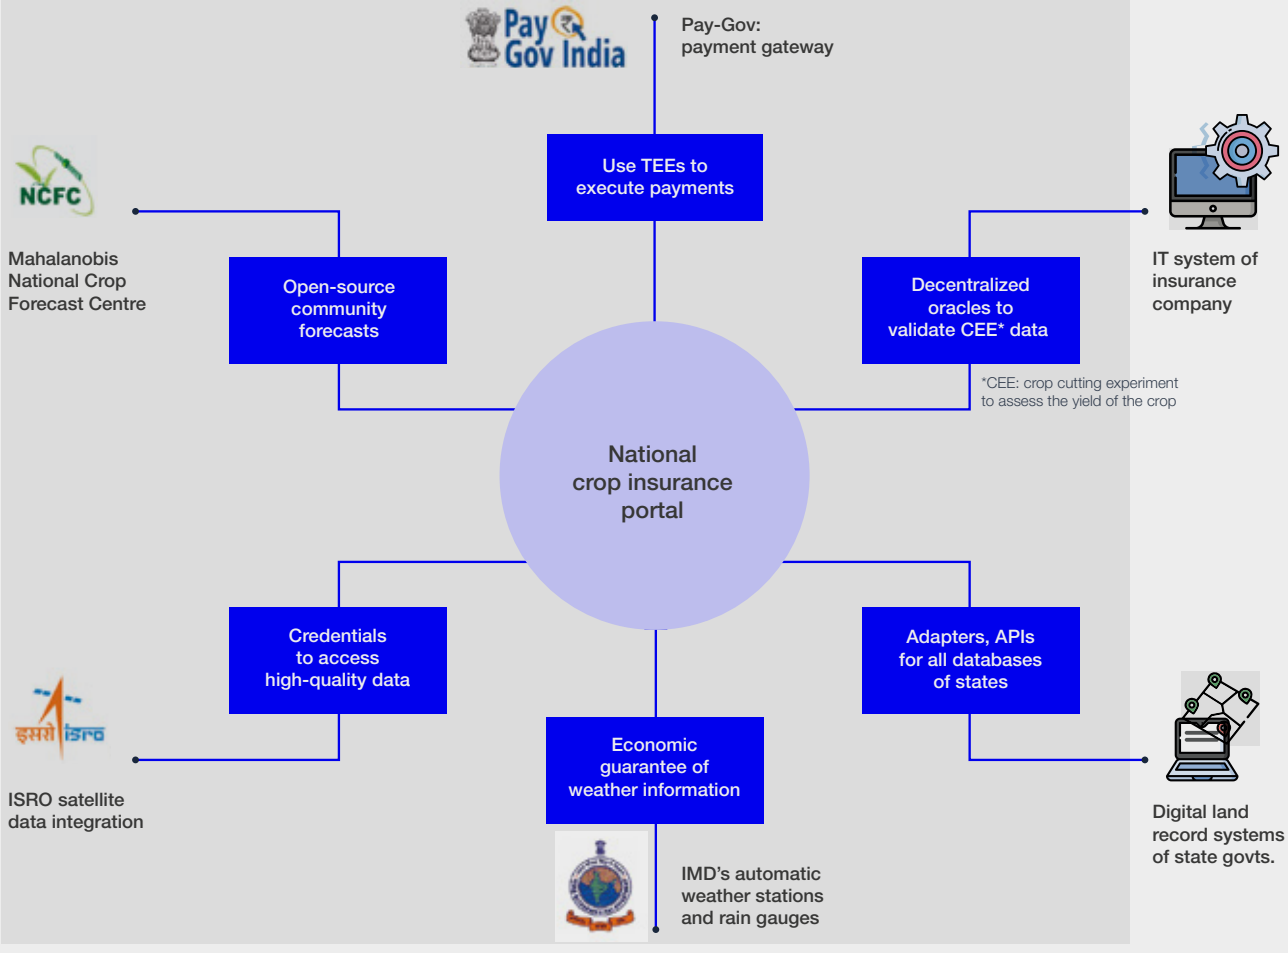

**Other potential use cases**

**Vehicle registration and monitoring for law enforcement**: The Government of India's Ministry of Road Transport and Highways standardized the registration processes of centralized applications (Vahan for vehicle registration and Sarathi for driving licences), which compile data from the state and national registers. Every registration application requires validation of vehicle and owner identification, address proof, owner credit score, insurance policy, invoice, taxation details etc. It may take up to 20 working days to validate such data, often using nine physical documents and four API services provided by public and private centralized systems.

**Customs processing and solving frauds/disputes regarding country of origin:** The global trade process involves extensive data and document sharing for compliance and clearance from both origin and destination country stakeholders. An end-to-end trade transaction may involve more than 28 different stakeholder systems. With the adverse impact of the COVID-19 pandemic, ongoing trade disputes and disruptive events such as Brexit, customs organizations regularly face frequent disputes over goods' country of origin, tax evasion and fraud.

Notes
1. By analysing mirror data on trade between China and Hong Kong, Fisman and Wei (2004) estimate that a 1% increase in taxes is associated with a 3% increase in fraud.
2. How Budget Counters 'Origin Fraud' in FTAs: https://www.thehindubusinessline.com/opinion/how-budget-counters-origin-fraud-in-ftas/article30794429.ece.

This paper attempts to suggest approaches that could be adopted and further developed to encourage interoperability between blockchain systems and legacy systems without making the claim that legacy systems need necessarily be replaced or that blockchain technology should unquestionably be employed. The approach to building blockchain middleware as the middle layer that can bring both the systems together will help in accelerating the positive effects of blockchain and distributed ledger technology where they are used and bring out the best of both worlds.

Smart contracts and other innovations in blockchain and distributed ledger technologies may significantly alter the way in which business transactions are currently undertaken. Data immutability, fault tolerance, censorship resistance and decentralization are elemental innovations that may lead to a fundamental change in how DLT is perceived by traditional organizations. However, legacy systems for data, communications and computations are the dominant tools for businesses and governments and are likely to be so for the foreseeable future. DLT-legacy interoperability furthers the potential of DLT and smart contracts by enabling DLT to engage with the real, physical world and apply those elemental innovations in physical-world use cases – as can be witnessed in a vast number of experiments happening worldwide (see box for other potential use cases).

DLT-legacy interoperability is a critical step towards unlocking DLT experiments outside the proof-of-concept pilot zone and bringing those to enterprise scale. The pathways suggested in this paper detail some possible approaches for building interoperability bridges to enable all relevant stakeholders to access the potential benefits of DLT and legacy systems and, in turn, generate innovative value creation in the economy. The social impact of DLT-legacy interoperability can also be very significant, whereby challenges of lack of trust and intermediary dependence in legacy systems can be overcome by innovations brought about by DLT. This paper hopes to kickstart a larger ecosystem effort towards unleashing these social and economic benefits by using the suggested pathways to build DLT-legacy interoperability. Policy-makers, government institutions and enterprises should evaluate these suggestions as per their operational environment and initiate building capabilities forexploiting employing DLT and smart contracts for their existing legacy systems without the need to replace them.

# Contributors

## Lead authors

**Sergey Nazarov**
Co-Founder, Chainlink
Chief Executive Officer, Chainlink Labs

**Punit Shukla**
Project Lead, Blockchain and Digital Assets
World Economic Forum

## Contributors

**Avery Erwin**
Head of Content, Chainlink Labs

**Amey Rajput**
Fellow, Blockchain and Digital Assets,
World Economic Forum, India

# Appendix: Reviewers

**Marcos Allende Lopez**
Tech Lead, LACChain
IT Specialist, Inter-American Development Bank

**Praphul Chandra**
Founder and Chief Executive Officer
KoineARTH

**Tas Dienes**
Ecosystem Support Program
Ethereum Foundation

**Arushi Goel**
Independent Consultant, UAE
Former Judge, Indian judiciary

**Nadia Hewett**
Project Lead, Blockchain and Digital Assets
World Economic Forum

**Purushottam Kaushik**
Head, Centre for Fourth Industrial Revolution India
World Economic Forum

**Dilip Krishnaswamy**
Vice-President (New Technology R&D)
Reliance Industries

**Ashley Lannquist**
Project Lead, Blockchain and Digital Assets
World Economic Forum

**Gordon Ian Myers**
Chief Counsel, Technology and Private Equity
International Finance Corporation
The World Bank Group

**Anirudh Rastogi**
Founder and Chief Executive Officer
Ikigai Law

**Jaideep Reddy**
Lead, Technology, Media and Communications
Nishith Desai Associates

**Ratul Roshan**
Associate
Ikigai Law

**Thibault Schrepel**
Faculty Affiliate at the CodeX Center (Stanford University)
Assistant Professor at Utrecht Law School
Invited Professor at Sciences Po Paris and University of Paris Sorbonne

**Nitin Sharma**
Founder, firstprinciples.vc
Founder, Incrypt Blockchain

**Sheila Warren**
Head of Data Policy, Blockchain and Digital Assets
Member of the Executive Committee
World Economic Forum

# Further reading


1. World Economic Forum, Blockchain Beyond the Hype, April 2018: http://www3.weforum.org/docs/48423_Whether_Blockchain_WP.pdf (link as of 25/11/20).
2. Enterprise Times, Enterprise Blockchain Adoption: 8 Reasons to Doubt, February 2020: https://www.enterprisetimes.co.uk/2020/02/24/enterprise-blockchain-adoption-8-reasons-to-doubt-part-i/ (link as of 25/11/20).
3. R3, Blockchain: 2020 Vision https://www.r3.com/wp-content/uploads/2020/01/blockchain_2020_vision_shaping_enterprise_blockchain_adoption.pdf (link as of 25/11/20).
4. Forbes, Blockchain's "'Troughs of Disillusionment'" Are Really the "'Trenches of Deployment'", January 2020: https://www.forbes.com/sites/richardgendalbrown/2020/01/09/blockchains-troughs-of-disillusionment-are-really-the-trenches-of-deployment/?sh=41907b955f65 (link as of 25/11/20).
5. Gartner Top 10 Strategic Technology Trends for 2020, October 2019: https://www.gartner.com/smarterwithgartner/gartner-top-10-strategic-technology-trends-for-2020/ (link as of 25/11/20).
6. FAO, Blockchain for Agriculture: Opportunities and Challenges, 2019: http://www.fao.org/3/CA2906EN/ca2906en.pdf (link as of 25/11/20).
7. Financial Times, Smart Insurance Helps Poor Farmers Cut Risk, December 2018: https://www.ft.com/content/3a8c7746-d886-11e8-aa22-36538487e3d0 (link as of 25/11/20).
8. World Economic Forum, Inclusive Deployment of Blockchain for Supply Chains: Part 6 – A Framework for Blockchain Interoperability, April 2020: https://www.weforum.org/whitepapersinclusive-deployment-of-blockchain-for-supply-chains-part-6-a-framework-for-blockchain-interoperability (link as of 25/11/20).
9. Hackernoon, Understanding the Gold Rush of Scalable and Validated Data Powered by Blockchain, March 2018: https://hackernoon.com/understanding-the-gold-rush-of-scalable-and-validated-data-powered-by-blockchain-and-decentralized-ee05db6b6a68 (link as of 25/11/20).
10. Forbes, Blockchains Are Verticalizing, So We Need Interoperability, March 2018: https://www.forbes.com/sites/adrianbridgwater/2018/02/07/blockchains-are-verticalizing-so-we-need-interoperability/?sh=263dbbd77ab9 (link as of 25/11/20).
11. Global Innovation Lab for Climate Finance, Blockchain Climate Risk Crop Insurance, September 2019: https://www.climatepolicyinitiative.org/wp-content/uploads/2020/08/Blockchain-Climate-Risk-Crop-Insurance_instrument-analysis.pdf (link as of 25/11/20).
12. Medium, Decentralised Oracles: A Comprehensive Overview, January 2019: https://medium.com/fabric-ventures/decentralised-oracles-a-comprehensive-overview-d3168b9a8841 (link as of 25/11/20).
13. R3, Blockchain Byte: R3 Research: https://www.finra.org/sites/default/files/2017_BC_Byte.pdf (link as of 25/11/20).
14. Forbes, The Five Ingredients of Blockchain Interoperability, February 2019: https://www.forbes.com/sites/richardgendalbrown/2020/02/13/the-five-ingredients-of-blockchain-interoperability/?sh=6df1b2e958a1 (link as of 25/11/20).
15. Jason Teutsch, On Decentralized Oracles for Data Availability, December 2017: https://people.cs.uchicago.edu/~teutsch/papers/decentralized_oracles.pdf (link as of 25/11/20).
16. Polkadot, An Introduction to Polkadot: https://polkadot.network/Polkadot-lightpaper.pdf (link as of 25/11/20).
17. Cosmos, What is Cosmos?: https://cosmos.network/intro (link as of 25/11/20).
18. Enterprise Ethereum Alliance, Crosschain Interoperability Use Case, June 2020: https://entethalliance.org/wp-content/uploads/2020/08/CIFT_Use_Case.pdf (link as of 25/11/20).
19. Baseline Protocol: https://docs.baseline-protocol.org/ (link as of 25/11/20).
20. Nick Szabo, Money, Blockchains and Social Scalability, February 2017: https://unenumerated.blogspot.com/2017/02/money-blockchains-and-social-scalability.html (link as of 25/11/20).
21. The European Union Blockchain Observatory and Forum, Scalability Interoperability and Sustainability of Blockchains, 2020.
22. Vitalik Buterin, SchellingCoin: A Minimal-Trust Universal Data Feed, March 2014: https://blog.ethereum.org/2014/03/28/schellingcoin-a-minimal-trust-universal-data-feed/ (link as of 25/11/20).
23. Michael J. Schallop, The IPR Paradox: Leveraging Intellectual Property Rights to Encourage Interoperability in the Network Computing Age, AIPLA Quarterly Journal 28, no. 3 (2000): 195; also, Michael A. Carrier, Innovation for the 21st Century: Harnessing the Power of Intellectual Property and Antitrust Law (Oxford University Press, 2009), and Jean Tirole, The Theory of Industrial Organization, (MIT Press, 1994): 405.

24. Richard M. Steuer, Standard Setting: Can Be Rife With Opportunities for Anticompetitive Activity," June 2016: https://www.mayerbrown.com/en/perspectives-events/publications/2011/06/standard-setting-past-present-and-future (link as of 25/11/20).

25. Jean Tirole, Normes et Propriété Intellectuelle: La vue d'un économiste, Lettre de l'Autorité de Régulation des Communications Electroniques et des Postes 51, 2006.

26. Bowen Liu, Pawel Szalachowski, A First Look into DeFi Oracles: https://arxiv.org/pdf/2005.04377.pdf (link as of 25/11/20).

# Endnotes

1. Gartner projects that banks will derive ~$1 billion of value using blockchain technologies by 2020 (Gartner's Top Strategic Technology Trends for 2021): https://www.gartner.com/smarterwithgartner/gartner-top-strategic-technology-trends-for-2021/ (link as of 24/11/20).
2. This is broken when a majority of nodes collude together to change the validation rules/outcome, which is highly unlikely to happen given the incentive structure of the underlying blockchain protocol – although some consensus protocols have instantaneous finality.
3. PMFBY Dashboard: https://pmfby.gov.in/ceo/dashboard (link as of 24/11/20).
4. See ConsenSys Quorum: https://consensys.net/quorum/; Github, ConsenSys Constellation: https://github.com/ConsenSys/constellation (links as of 24/11/20).
5. Source: ISO/TC 307 for standard definitions.
6. World Economic Forum, *Inclusive Deployment of Blockchain for Supply Chains: Part 6 – A Framework for Blockchain Interoperability*, April 2020: https://www.weforum.org/whitepapers/inclusive-deployment-of-blockchain-for-supply-chains-part-6-a-framework-for-blockchain-interoperability (link as of 24/11/20).
7. The concept of abstraction is used extensively in computer science to focus on elements of larger importance and ignore other temporal details. For example, the use of data types abstracts away all of the internal details of how a variable is stored and processed in a computational system and allows the system to interact with any kind of data as an abstracted data entity.
8. Jean Tirole, Normes et Propriété Intellectuelle: La vue d'un économiste, Lettre de l'Autorité de Régulation des Communications Electroniques et des Postes 51, 2006.

WORLD ECONOMIC FORUM

COMMITTED TO IMPROVING THE STATE OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland

Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org