Techruption Consortium Blockchain – what it takes to run a blockchain together

Oskar van Deventer
TNO
Den Haag, Netherlands
oskar.vandeventer@tno.nl

Frank Berkers
TNO

Mischa Vos
Rabobank

André Zandee
APG

Tom Vreuls
APG

Laurens van Piggelen
APG

Alexander Blom
Bloqzone

Bas Heeringa
BSSC

Saïd Akdim
Kamer van Koophandel

Paul van Helvoort
CZ

Leon van de Weem
Zuyderland Medisch Centrum

Douwe van de Ruit
KPN

ABSTRACT
This paper presents initial results of the Techruption Consortium Blockchain experiment. The purpose of the experiment is to learn what it takes to run a permissioned consortium blockchain infrastructure together, not only from a technical perspective, but also governance and business model. The experiment turned out to be surprisingly complex, running into buggy open-source software, extensive firewall and connectivity issues, a complex legal context, a plethora of governance issues, many business model alternatives, and an ever-present human resource limitation. Based on our experiences, we conclude that instead of developing dedicated technical infrastructure, governance and business models for each blockchain application individually, there is a need for a shared blockchain infrastructure with basic governance and business models to spur further innovation in blockchain applications and enabling technologies.

Author Keywords
Blockchain; consortium blockchain; permissioned blockchain, governance, operational/technical, business model, experiment

ACM Classification Keywords
H.4.m. Information systems applications: Miscellaneous

van Deventer, Oskar; Berkers, Frank; Vos, Mischa; Zandee, André; Vreuls, Tom; van Piggelen, Laurens; Blom, Alexander; Heeringa, Bas; Akdim, Saïd; van Helvoort, Paul; van de Weem, Leon; van de Ruit, Douwe (2018): Techruption Consortium Blockchain – what it takes to run a blockchain together. In: W. Prinz & P. Hoschka (Eds.), Proceedings of the 1st ERCIM Blockchain Workshop 2018, Reports of the European Society for Socially Embedded Technologies (ISSN 2510-2591),
DOI: 10.18420/blockchain2018_06

INTRODUCTION: BLOCKCHAIN
Blockchain has been receiving a lot of industry attention for the last few years. Inspired by the illustrious Bitcoin [1] system, initiated by the illusive Satoshi Nakamoto, many other blockchain technologies have been developed as well as numerous applications that rely on a blockchain. Some blockchain technologies are relatively general-purpose, like Ethereum [2] for enforcing smart contracts and Hyperledger Fabric [3] for running chaincode, whereas other blockchain technologies have more specific purposes, like Sovrin [4] for identity transactions, Ripple [5] for financial transactions, and BigchainDB [6] for long-term data storage.

Siegel [7] explains blockchain as “A blockchain is a shared ledger that everyone trusts to be accurate forever”. A ledger is a record of transactions. Shared means that there is a single ledger that is the same for all participants. This combination is a unique selling point of blockchains, alleviating the efforts for synchronization between ledgers of individual participating organizations or individuals, and hence reducing transaction cost and bureaucracy. Trust is the keyword. Participants are no longer required to trust a single organisation for the contents of the shared ledger, but they trust a blockchain business ecosystem instead, where no single party has the power to make unauthorized changes to recorded transactions.

A blockchain infrastructure consists of nodes (servers) that are run by organizations and individuals that have an incentive to run a part of the infrastructure. Different terms are used for this business role like miner, validator, steward and server, depending of the technology that is used. We shall use the generic term “blockchain service provider”. A blockchain service provider enables others to submit transactions to the blockchain and to read transactions from the blockchain. The group of blockchain service providers validates and confirms transactions to maintain a single consistent ledger.

Mark Peplow et al [8] distinguish two dimensions to characterize blockchains: permissionless-permissioned and public-private. The first dimension is about who can act as blockchain service provider. In a permissionless ("unpermissioned") blockchain, anyone can validate and confirm transactions and consensus algorithms like proof-of-work and proof-of-stake are used to keep the blockchain consistent. In a permissioned blockchain, only identified participants can validate and confirm transactions, and some type of "byzantine-fault-tolerant" voting mechanism is used to keep the blockchain consistent [18]. The second dimension is about who can access and use the blockchain. In a public blockchain, anyone can perform blockchain transactions. In a private blockchain, only identified participants can perform blockchain transactions.

Based on these two dimensions, three types of blockchains may be distinguished [8], see Figure 1.

1) Cryptocurrency blockchain
2) Private blockchain
3) Public consortium blockchain

	Permissionless: **Anyone can validate and confirm transactions**	**Permissioned:** **Only participants can validate and confirm transactions**
Public: **Anyone can perform blockchain transactions**	1) "Cryptocurrency blockchain"	3) "Public consortium blockchain"
Private: **Only participants can perform blockchain transactions**	N/A	2) "Private blockchain"

Figure 1. Categorization of blockchain types

Cryptocurrency blockchains like Public Bitcoin and Public Ethereum have the benefit that they are up and running, and they have a proven past performance. They are public, i.e. visible and accessible to anyone. They are permissionless, so nobody can prevent one from participating. This is why start-ups typically use a cryptocurrency blockchain for their to-be-trusted core applications. However, cryptocurrencies are volatile and their blockchains are frequently forking, which is how these blockchains are governed. Escalating transactions fees and transaction confirmation times have become an issue for applications and exchanges [19]. Regular forking implies that an industry sector would continuously need to coordinate and resolve which fork to use.

Private blockchains, e.g. based on Hyperledger Fabric or Ethereum Enterprise technology, have the benefit that partners do not need to rely on others for blockchain access. Instead, the consortium partners that have agreed on an application-specific bureaucracy-reduction solution are the participants in their own joint private blockchain. Several industry sectors are already developing their own private blockchain solutions, in many cases supported by an American tech giant. Private blockchains also have disadvantages. As we are learning from Techruption Consortium Blockchain (see below), it is complex and costly to run a blockchain network together. Outsourced blockchain operations run the risks of technology and vendor lock-ins. Also, the appearance of a plethora of application-specific micro blockchains impedes trans-sectoral innovation, like combining blockchain solutions on finance and logistics, or energy and identity.

A public (permissioned) consortium blockchain facility, which may also be based on Hyperledger Fabric or Ethereum Enterprise technologies, is run by a consortium of participants for whom trust is their core business and who have the expertise to run such a facility as efficiently and reliably as possible, e.g. banks and telecom operators. A public consortium blockchain may have lower operational cost, and hence lower transaction fees, than (networks of) private blockchains. They may resolve several of the above-mentioned issues with private blockchains. Also they allow industry sectors to focus on their core business, which running blockchains is not. Whereas there are already some public consortium blockchains in existence (e.g. Sovrin for identity solutions, and Interplanetary Database for data storage), the concept is still far from being a proven solution.

TECHNOLOGICAL TURBULENCE

As illustrated above, blockchain technology is in practice a collection of technologies, both complementary and mutually exclusive ones. Moreover, this class of technologies is still heavily in development, with new alternatives and trials popping up virtually on a daily basis. This characterises the turbulence associated with the market adaptation phase [16] and brings forward high levels of uncertainty to those seeking to utilize the immense promises of these technologies. Two major strategies are available: wait until the market is stabilized or explore how the characteristics of the technology can be put to benefit. The infrastructural or platform character of the technology implies that the blockchain technology can be used generically to enable applications. This puts forward the challenge to understand application requirements and implications for this infrastructure [15] and vice versa – in order to be able to recruit a critical mass of users of the platform. The huge difference between the number of ideas and the number of active blockchain applications suggests that also in the area of blockchain enabled applications a phase of exploration is pertinent. However, if we can manage the technical uncertainty and provide a reliable infrastructure to potential blockchain enabled application, this is likely to lead to a flux of innovation initiatives. For many, mostly incumbent organisations both these uncertainties as well as the open and fully decentralized character drive their need for experimentation, but in a more controlled environment. The inherent characteristic of blockchain technology to involve multiple parties drives the desire to do experimentation in a joint controlled environment.

TECHRUPTION CONSORTIUM BLOCKCHAIN

Techruption Consortium Blockchain is a project within the 15+ partner Techruption [9] program. The program aims to jointly develop use cases on a.o. blockchain. The project partner are a bank, a pension fund, a health insurer, a hospital, Dutch chamber of Commerce, a telco, a start-up and a research institute. The rationale of this project is that any blockchain initiative will sooner or later run into a make-or-buy decision with respect to the infrastructure on which the blockchain application runs. Will the blockchain application run on a blockchain infrastructure that is newly-created by the initiative itself? Will the blockchain application run on a blockchain that is run by third parties? What type of blockchain should the application run on? Etcetera.

When the project started spring 2017, the partners (see author list) realized that more research was needed to provide substantial insight in how to make the above-mentioned make-or-buy decision. How difficult is it to run a blockchain from the technical perspective? How much does it cost? What is the business model? What would the governance look like? What components can be outsourced/bought? How mature is the technology? What are the risks. Etcetera.

The partners are participating in the project for a multitude of reasons, which differ per partner.

- Getting experience in developing private blockchains, and using that experience for the own industry sector.
- Developing a blockchain platform infrastructure for research use and research projects.
- Developing blockchain applications, and using the developed infrastructure for technical and market testing.
- Understanding governance requirements.
- Executing business simulations.
- Working toward a professionally-run public Dutch blockchain facility, possibly including neighbouring countries.
- Networking and collaboration opportunity.

Our main research questions for this experiment was “What does it take to run a consortium blockchain together, in setting up and managing technology? What governance model and business model is suitable for managing the consortium blockchain?”. We found that an experiment – as a study for a potential exploitation - with either one of these three dimensions would fall short. Running a consortium blockchain requires a coordination with respect to e.g. versioning, sharing of node-IPs, testing and monitoring and consequently some form of governance would be required. Coordinated decisions on e.g. technology, number of nodes etc. affects performance and functionality of the consortium and consequently potential value of the infrastructure for intended users. This implies that technology and governance are relevant for the exploitation and vice versa. Therefore we chose to experiment with technology, governance and business modelling intertwined.

RESEARCH APPROACH

We followed the “Groeifabriek” approach as an innovation management approach [10] to guide us from ideation to our current stage. The chosen research methodology is action research, which is a structured form of learning by doing [11]. We approached our research questions with respect to technology, governance and business as a practical experiment in which we try to establish the required technology and define the governance and business approach for exploiting the consortium infrastructure. In this experiment we tried to apply the chosen governance principles to the project as if we were a consortium actually exploiting the infrastructure. In order to capture the consensus view on these aspects, we kept to a contribution driven blueprint document that requires contributions to be approved by decision meetings.

In the ideation stage, the general direction of the project was decided. In the exploration stage, we developed an initial business model and we set up an initial governance structure. In the experimentation stage, we experimented with the technology, governance and business models, documenting our joint consensus vision in a blueprint document [12].

The remainder of this paper discusses the learning experiences from the technical infrastructure experiments, governance design and experiences, business modelling and phasing, and envisioned next steps.

TECHNICAL INFRASTRUCTURE EXPERIMENTS

The technical infrastructure experiments were kicked off by inviting Accenture, technical partner of Techruption, to guide us through the technology selection and instantiation process.

Accenture used their Blockchain Vendor Assessment Framework to show the different options in a structured way. The process involved deciding about the layer in which the project is active (applications and solutions, platforms, base protocols & infrastructure), the functionality of the infrastructure (transaction processing and data storage, basic distributed execution platforms, advanced distributed execution products), as well as practicalities like available technology expertise and architecture design.

Quorum [13] was chosen as blockchain technology. Quorum is a permissioned variant of Ethereum technology, and compatible with Ethereum at the application level. Quorum 1.2 was the latest version when the experiments started. Each of the Quorum node instances is running on a software stack with Docker and Linux to provide flexibility for instantiation, moving and future upgrading of nodes, see Figure 2. The software stack is maintained on a joint GitLab repository to enable proper joint version control.

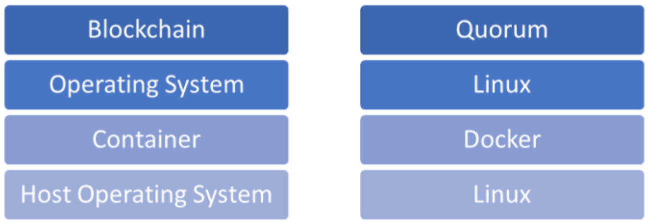

Figure 2: Software stack.

Different partners use different platforms to run the software stack: own cloud infrastructure, third-party cloud infrastructure and even a RaspberryPi. The blockchain was initiated during a one-day workshop. A genesis block was created, IP addresses and Quorum enode IDs were exchanged, firewalls were opened, connectivity was tested and the blockchain was started.

Many problems arose during the six months that the technical infrastructure is running. In many cases, it was hard to make proper technical diagnoses or find robust solutions, e.g. relating to bugs in the Quorum releases or IT settings or our own organizations, as there is a lack of central monitoring tools (who will trigger actions when some nodes are not syncing well?), whereas system logs and port scans turned out little information. Other problems were more of an organizational nature, e.g. getting the right people to develop contributions.

It took a full five days to get all five initial nodes connected and synchronized. We still do not know why this took so long. Many parameters need to be configured in the software stack, and we are only gradually learning their impact. At many times, nodes went down for unexplainable reasons, leading to speculation about the robustness/bugginess of the Quorum software. At many times, the network was diagnosed to be less than a full mesh. Lots of work went into configuring and reconfiguring firewalls, as many nodes were restarted from scratch from a different IP address. One partner had a dedicated cloud infrastructure for this type of systems, but the administrator refused access for unclear/bureaucratic reasons. One partner has a system with a multitude of firewalls, where each minor firewall reconfiguration requires a call to a helpdesk and a one-day execution time. One partner has an "intrusion prevention system" that intercepts, decrypts and re-encrypts all traffic, leading to major SSL certificate issues. Two of the nodes permanently crashed when the blockchain outgrew the assigned memory allocation. During the experiment, a new Quorum 2.0 version was released. One partner was unable to make the upgrade, whereas some other partners were unable to maintain the deprecated version. We were unable to smoothly migrate/fork the state of the initial network to Quorum 2.0, so we decided launch an independent Quorum 2.0 network.

The good news is that there were no major issues at the application level. We have successfully deployed and interacted with a multitude of Ethereum/Solidity smart contracts, including a "hello-world" smart contract, a VoIP communication-management application by one of the partners, and a self-sovereign-identity application from a neighbouring project of the Techruption Blockchain program.

As can be derived from the above, many of the technical challenges we encountered were of a generic nature, e.g. firewall configuration, access to skills, buggy software and the coordination required to diagnose it. This suggests that setting up consortium blockchains could benefit from proper design of governance and allocation of personnel for setting up, testing, debugging and accepting the technology.

GOVERNANCE DESIGN AND EXPERIENCES

The governance design was kicked off using a governance framework developed by TOBLOCKCHAIN, a blockchain start-up and partner of Techruption. Three questions are central in this framework.

- What issues can the decisions be about?
- How are decisions made? (decision process)
- Who participates in the decision process?

The list of potential issues governed by the decision process is a long one, including technical choices for the software stack, version control of the used third-party open-source software, technical requirements on connectivity and firewalls, monitoring and maintenance of key performance parameters, division of cost and revenues, procedures for onboarding new partners and new customers, procedures for offboarding, business model and phasing, and of course the decision process itself.

The governance of the *project* was split in an informal process and a decision process, see Figure 3. During the informal process (workshops, etcetera) opinions are formed and consensus is sought. Volunteers make written proposals and change request based on this. The decision process is centred around a blueprint document, that is updated at every decision meeting (semi-weekly conference call), based on decisions on the provided inputs, see Figure 4.

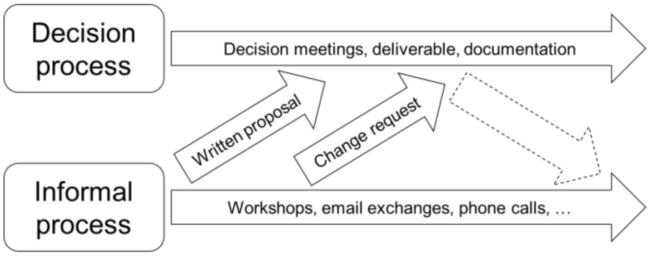

Figure 3: Governance of the project

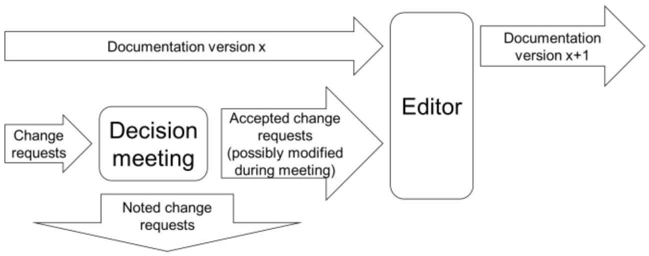

Figure 4: Handling change requests at a decision meeting

As one could expect, the practice was a bit more complicated. The initial version of the governance process, as we designed for the *future ecosystem*, requires full consensus between the founding fathers of the project. However, we never had a full set of representatives present at our decision meetings due to conflicting appointments, illness and so on. Also, one partner withdrew during the project, while two others joined. A more practical approach was to achieve consensus between those present at the meeting, and assuring that enough people are present. This shows that the governance framework is dynamic.

Also the volunteer-based contribution-driven approach has its limitations. Having multiple authors has its quality impacts, including variations in writing style and terminology. Moreover, not all partners could contribute equally, so some partners carried a larger contribution load than others. Still, a strong point of the chosen approach is that all partners have a stake in the resulting blueprint, as it was developed by the partners themselves and contributions were included by consensus.

The experience suggests that the governance model needs to be adaptable to the (increasing) complexity of the situation and yet be pragmatic. As can be observed in the governance of many foundations, working groups lead by champions that focus and take responsibility on certain aspects could have contributed to a more smooth advancing of topics – as opposed to the all-contribute all-decide approach.

BUSINESS MODELLING AND PHASING

Similar to the governance and technology tracks, this topic was initiated with an informal and interactive session based on the business model canvas [14] led by the Groeifabriek.

In order to further facilitate the scoping choice of the consortium, TNO prepared a so called strategic options model, see Figure 5. This model combines the phases of innovation with platform business model theory [15]. The latter essentially distinguishes a platform or infrastructure on which multiples sides interact (also referred to as multi-sided market). In our case these sides are represented by the applications that require a blockchain ("demand") and the providers of components that jointly make up the blockchain infrastructure ("supply"). These sides theoretically have network effects as blockchain applications only make sense if the underlying infrastructure is sufficiently large in number of nodes and has sufficient support. On the other hand, contributing to a blockchain infrastructure only makes sense if sufficient applications utilize this infrastructure. The innovation management perspective distinguishes exploration (or experimentation) from exploitation [16]

In the exploration phase, new technologies are tried and tested and eventually, if considered feasible applied for exploitation. This means that the technology is actually used in business. The combination of the two leads to distinguish exploration and exploitation in both applications as well as platform. Moreover, if considered more closely, it would make sense to exploit an infrastructure that is specifically suited to support the experimentation with blockchain applications in order to spur the innovation of blockchain applications (A2 in Figure 5). After all, value lies in the actual use of applications (A3 in Figure 5).

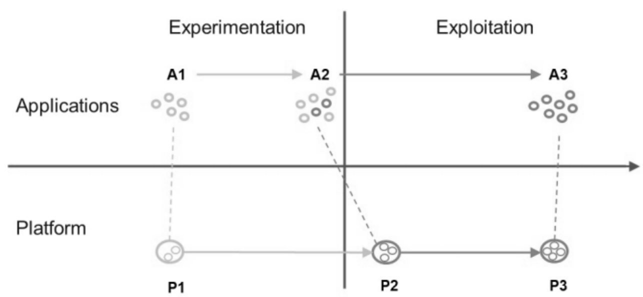

Figure 5: TCB Strategic Options Model

We named this platform "P2", a 'professionally ran platform for the experimentation with blockchain applications', to distinguish the phase of the platform in between "P1", where the platform itself is experimental (current phase) and "P3", where the platform is suited for real-life blockchain applications that rely on existence of a well-managed infrastructure. This distinction is similar to the testnets for public blockchains. The consortium unanimously opted for the P2 scope as a target for the consortium blockchain. This is all further documented in [12]. The P2 platform has value for its founding fathers who have plenty of blockchain application ideas as well as start-ups interested in developing blockchain or complementary applications.

Later on, based on the implied analogy with systems implementation procedures that typically distinguish development, test, acceptance and production [17] it was established that "P3" would be the 'acceptance test platform' which closely resembles the production environment and "P4" is the production environment. This distinction did not affect the choice for P2.

At that time, it was unclear whether the consortium was considering the platform to evolve from P1 to P2 etc. or whether P1, as an explorative infrastructure, could exist besides P2. Based on the current pace of development in blockchain technologies, the consortium still has a need to explore such technologies, hence chose to maintain a P1 instance besides P2.

In order to specify the value proposition [20] of the consortium blockchain in more detail another informal session was devoted to identify the service elements that TCB provides to its users. It was decided that TCB would not support the application developers' development phase. This is something that developers typically do "offline".

At this point TCB neither provides support for acceptance testing or production as the specs for such support are not clear yet. Thus TCB focuses on the 'test phase' for application developers. The support can be split into two categories: during test and prior to test. The services provided for testing are fairly basic and include the necessary APIs and monitoring and alerting. No 24/7 helpdesk will be operational, but support will be provided at best effort through the TCB community and an escalation/routing mechanism.

The services prior to testing are mainly informational services that the developer needs to prepare for the testing. These include: configuration information; platform status and performance information; "Service Level Agreement" (what a developer can expect); a roadmap of development plans for additional functionalities; a manual that describes procedures for deployment, testing and decommitting; deployment automation software; release management systems; optional generic functionality (e.g. identity management) that can be included in the applications (e.g. as libraries); procedures for proposals from the developers to the infrastructure; cost and performance information.

Currently these services are targeted at the parties and their ventures that are part of the Techruption community. This puts BSSC, as the host organization for the Techruption community, in a key-role to adopt and orchestrate the further development of the TCB. The Techruption community is open for participation by third parties. At this point, and based on a developed roadmap for technical, organizational and ecosystem development, TCB estimates to need around 6-9 months to fully achieve this "P2" stage.

NEXT STEPS

With the first project phase completed, the partners are now at the stage of committing to the next phase, which includes a next level of professionalization. More clarity needs to be obtained on the cost of running this blockchain ecosystem, the value of the ecosystems and transactions to its customers, the acquisition of such customers, tariffing models and further professionalizing the governance. There is also work to be done at the technical level, including improving (dockerised) template deployment to minimize the faults and standardise the configuration, network monitoring to identify failing nodes, strict rules and rigid mitigations on version management, handling of deprecations in smart contract languages, and implementing security measures.

At the time of writing this paper, there are still a significant number of open issues to be addressed for the exploitation phase. Should we engage in a utility service? How can we benefit as a business? Should this all be governed by an umbrella organisation? Should we go commercial or non-profit? How open should the platform service be? Would it run only private blockchains or also public blockchains? Should permissionless blockchains also be considered? Although fundamental for the organisation's future in actually using blockchain technology in their business, these issues do not disqualify the need for "P2". On the contrary, further development and exploration is needed to answer these questions.

REFLECTION AND CONCLUSIONS

The experiment was evaluated using brainwriting and face to face discussion and focused on identifying aspects to maintain and things to improve in the collaboration. Aspects that emerged (no prior structure was given) include strategy, scope and output as well as operational organisational aspects. Key insights beyond a unanimous agreement on willingness to continue were that the group was perceived as multidisciplinary, open minded, constructive and very knowledgeable. Remarkably so, since participants were not selected on a specific profile. The group agreed that this was definitely an aspect to cherish. On a more critical note the group concluded that more emphasis should be on actually iteratively developing and further scaling the infrastructure and other achievements valuable for users of the infrastructure. This, rather than emphasis on, but explicitly not replacement of, analytical discussion and documentation. A more formal project plan was deemed instrumental. The group considered this now to be a natural moment for the shift from an informal, open and explorative experimentation phase to more structure, although some of the participants clearly desired that some time ago.

After nine months of action research, we conclude that setting up and running a consortium blockchain together is much more complex than we had anticipated. The group deliberately chose not to hire (turn-key) solution providers for either of three areas technology, governance and business in order to learn in practice. Consequently learnings include new aspects and aspects known (elsewhere). Overall it became clear that setting-up and running a consortium blockchain can be considered a business in its own right. We learned how much the software is still developing, that there are lots of things to be configured in the software stack and that it is hard to technically maintain a stable-running blockchain. Also the governance turned out to be complex, so many issues for which procedures would need to be developed, and already running into scale issues with only a handful of partners. Moreover, the experiment clearly illustrated how much technology, business and governance of a blockchain infrastructure at this phase of technological turbulence are intertwined. We believe that neither one of these aspects can be meaningfully developed in isolation. Group learning undeniably includes phases of getting to know each other and developing a common language. This feels unproductive in phases, however the different perspectives and backgrounds eventually added to the level of understanding. Thus, on top of the basic blockchain infrastructure and the governance and business principles documented in the blueprint, we consider the joint team as a fourth asset put forward by this experiment.

Together, we have developed a roadmap for which we have just completed the experimental “P1” phase, starting further professionalisation with the “P2” phase soon.

We recognize that many of the challenges that we have run into are independent of the number of blockchain instances that we are running. The governance processes are not much harder for running multiple blockchain instances, compared to just one. The same goes for technologies and business models. So it makes a lot of sense to run multiple blockchain instances and technologies in parallel, instead of developing dedicated technical infrastructure, governance and business models for each blockchain instance individually.

ACKNOWLEDGMENTS
This work was partly funded by TKI-toeslag.

REFERENCES

1. Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. https://bitcoin.org/bitcoin.pdf
2. Vitalik Buterin. 2013. Ethereum White Paper: a next generation smart contract & decentralized application platform. http://www.the-blockchain.com/docs/Ethereum_white_paper-a_next_generation_smart_contract_and_decentralized_application_platform-vitalik-buterin.pdf
3. Hyperledger Fabric. https://www.hyperledger.org/projects/fabric
4. Sovrin. For self-sovereign identity and decentralized trust. https://sovrin.org/
5. Ripple. One Frictionless Experience To Send Money Globally. https://ripple.com/
6. BigchainDB. The blockchain database. https://www.bigchaindb.com/
7. David Siegel. 2017. What is this Blockchain Thing?, https://medium.com/startup-grind/what-is-this-blockchain-thing-a5d2abb99297
8. Mark Peplow, et al. 2016. Distributed Ledger Technology: beyond block chain. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/492972/gs-16-1-distributed-ledger-technology.pdf
9. Techruption. 2016-… https://www.techruption.org/
10. GroeiFabriek. https://groeifabriek.com
11. Mårtensson, Pär, and Allen S. Lee. "Dialogical action research at omega corporation." MIS Quarterly (2004): 507-536.
12. Oskar van Deventer et al. Techruption Consortium Blockchain – blueprint of a permissioned blockchain. March 2018. Available through the first author.
13. J.P.Morgan. Quorum, advancing blockchain technology. https://www.jpmorgan.com/global/Quorum
14. Osterwalder, Alexander, and Yves Pigneur. Business model generation: a handbook for visionaries, game changers, and challengers. John Wiley & Sons, 2010.
15. Eisenmann, Thomas, Geoffrey Parker, and Marshall W. Van Alstyne. "Strategies for two-sided markets." Harvard business review 84.10 (2006): 92.
16. Ortt, J. Roland, and Jan PL Schoormans. "The pattern of development and diffusion of breakthrough communication technologies." European Journal of Innovation Management 7.4 (2004): 292-302.
17. DTAP, https://en.wikipedia.org/wiki/Development,_testing,_acceptance_and_production
18. Shehar Bano et al. SoK: Consensus in the Age of Blockchains. https://arxiv.org/abs/1711.03936v2. 2017.
19. Joseph Young, “Ethereum’s high fees have become an issue for applications and exchanges”, https://btcmanager.com/ethereums-high-fees-issue-for-applications-and-exchanges/, January 2018
20. Radziwill, Nicole. "Value Proposition Design." The Quality Management Journal 22.1 (2015): 61.