

## Open Call 2

### Decentralized data ecosystem for the Open Blockchain for Asset Disposition Alliance Deliverable 3: Experiment Results and Final Report

Authors	Leandro Navarro, Javier Cano, Adrián Manco, David Franquesa, Pedro Vílchez, Felix Freitag, Roc Meseguer, UPC; Rohi Sukhia, OBADA; Ronald S. Lembke, UNR.
Due Date	20210730
Submission Date	20210806
Keywords	Testbed, circular economy, digital devices, distributed ledger technology

## Deliverable 3: Part I

### Analysis, results, and wider impact

1.	Abstract .....	1
2.	Project Vision.....	1
3.	Details on participants (both EU and US) .....	1
4.	Results .....	2
4.1	Context: devices and their data .....	2
4.2	Context: safety properties .....	4
4.3	Data of interest in common .....	6
4.3.1	Devices.....	7
4.3.2	Proofs.....	8
4.3.3	Fees.....	8
4.3.4	Documents.....	8
4.3.5	Reports .....	8
4.4	Architecture .....	8
4.4	Experimentation .....	10
4.5	API.....	10
4.6	Testing .....	12
4.7	Experimental performance analysis .....	12
5.	Discussion and Analysis on Results.....	20
6.	Present and Foreseen TRL .....	22
7.	Exploitation, Dissemination and Communication Status .....	23
8.	Impacts .....	25
9.	Conclusion and Future Work .....	26
10.	References .....	27
11.	Glossary .....	28
12.	Workplan Progress and Travel Details.....	31
13.	Funds Utilisation Report .....	34

---



## 1. Abstract

More devices are sold every year than human beings are living on Earth. Decarbonisation is a must to tackle the environmental crisis and comply with the global warming objective of the Paris Agreement.

Therefore, we need to create effective market ecosystems capable of reusing ICT devices, extending their lifespan for new uses, instead of always manufacturing new ones, through reuse, repair, and ensure final recycling in a sustainable way.

In collaboration with the Obada.io organization and the University of Nevada-Reno in USA, the NGI eReuse-Ledger testbed is a permissioned distributed ledger to support experimentation about device traceability.

*Changes: no changes over D2, changed over D1. (97 words)*

## 2. Project Vision

We need effective circular market ecosystems capable of reusing ICT devices, extending their lifespan for new uses, instead of always manufacturing new ones, through reuse, repair, and ensure final recycling in a sustainable way. The collaboration takes place between the industrial experience of OBADA, the knowledge about the reverse supply chain of electronics at the University of Nevada-Reno in USA with the NGI eReuse-Ledger testbed in Europe. It offers an experimental distributed ledger for traceability that does not store details but only proofs (hashes) and economic deposits for devices as incentives. It is privacy-preserving while ensuring the trust, verifiability, irreversibility, tamper proof required safety properties for environmental accountability. The experimental results provide insights about performance and scalability as well as test the design of the operations and parameters of public API to record key operations.

*Changes: no changes over D2, changed over D1.*

## 3. Details on participants (both EU and US)

US:

Rohi Sukhia, is director of Obada LLC, the Open Blockchain for Asset Disposition Architecture, and CEO of Tradeloop, an EE graduate of Cornell University and Intel Corp. veteran. Expert member of ANSI committee in ISO TC 307. Role: coordination of software pilot and experimentation. He is supported by two developers (Andrii, Akshay) who are preparing the experiments and a prototype system from the USA side.



Ronald S. Lembke, is Associate Professor of Supply Chain Management, Chair of the Managerial Sciences Department, College of Business, University of Nevada, Reno. In the area of Reverse Logistics, he is Chair of the Standards Committee of the Reverse Logistics Association, and Chair of the RLA SQRL Code project. He sits in the Board of Advisors of the Reverse Logistics and Sustainability Council. Role: research and leadership on reverse-supply business models, coordination of standards documentation.

EU:

Leandro Navarro is Professor at the Department of Computer Architecture of the Technical University of Catalonia (UPC). Leandro has participated and managed the participation of the Distributed Systems research group (DSG) in several EC funded projects. Principal investigator in eReuse related projects NGI-Policy-in-Practice (2020-2021), DLT4EU (2020). He initiated in 2013 research about the circular economy of digital devices and digital ledger technology applied to traceability and accountability of digital devices and impact assessment. He is expert of the UN ITU-T SG5: Environment, climate change and circular economy, and coauthor of Recommendation ITU-T L.1024. Role: project coordination of the EU side. Supported by a team of researchers and developers at UPC.

The UPC team includes additional members: David Franquesa, a PhD student in his last year, working in collaboration with the USA partners to coordinate and plan experiments and testing. Pedro Vílchez, in charge of the server infrastructure maintenance and devops. Roc Meseguer and Felix Freitag, associate professors supporting the research. Javier Cano and Adrian Manco, developed several parts of testbed enhancements and supporting experiments on the testbed side.

*Changes: more details about participants.*

## 4. Results

### 4.1 Context: devices and their data

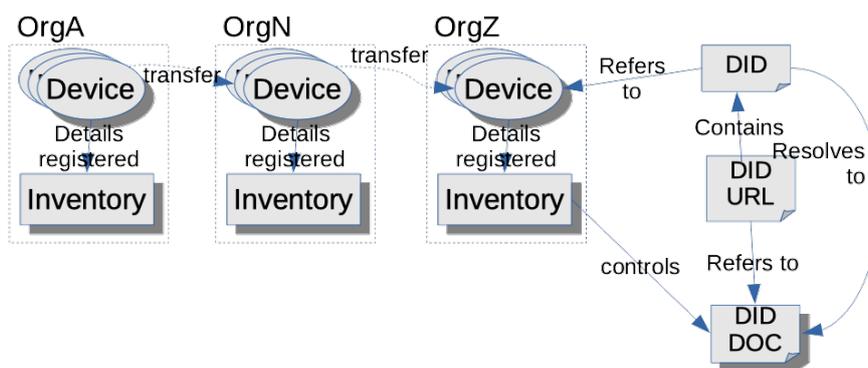
In a circular economy, the lifecycle of digital devices can be as follows: after the use of raw and secondary materials and parts, devices are assembled at factories and sold by brands. These devices have unique identifiers (serial numbers) that may come linked or labelled with details (information sheets) about their composition, characteristics, instructions for maintenance, repair, even recycling. In organisations, the details about computer devices are usually recorded in an inventory database, with associated information about insurance, maintenance, and accounting. Devices depreciate over time for a certain period. The end of one use cycle, when a device is no longer fit for its initial purpose, while it still has value for the owner organisation and maintenance, may create an opportunity for internal reuse for another less demanding purpose in the same organisation. When a device does not meet the needs of that organisation, or is too costly to be maintained or cannot be repaired, that marks



the end-of-use in that organisation. The device can be disposed of. However, disposed of devices in one organisation can be a resource for other users in the second-hand market, through refurbishment for device reuse, alternatively scavenging for parts reuse, or at least materials reuse from recycling, avoiding the highly pollutant burning or dumped in a landfill.

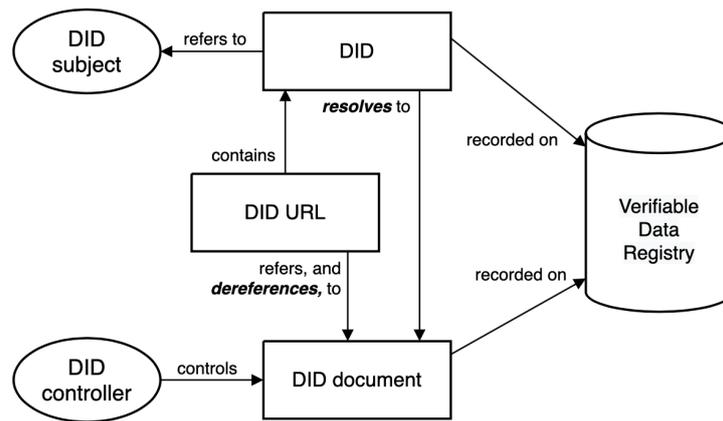
In summary, devices can be repaired, upgraded, transferred internally to a new use, or disposed to be transferred (sold, donated) to a new owner, dismantled for parts, or recycled to recover secondary raw materials or dumped in a landfill. Keeping an updated digital record of the history and status of a device and its parts is useful to facilitate circular processes, for the management of devices (keep track, particularly in volume), and for the accountability and verifiability [Küsters2010] of these processes, that have to do with compliance with requirements and expectations about business, environmental, social or economic aspects.

Digital support systems, like the eReuse ecosystem of tools have the structure in Figure 1. Multiple organisations (*OrgA..OrgZ*) have multiple devices ( $ID_1..ID_X$ ), with all details stored and updated as digital data in their organisational *inventory* system. Devices can be tagged with digital identifier codes using unique physical ID tags to facilitate identification for tracking and handling of these assets during the usage cycle, maintenance, end-of-use phase before final disposal. It is a common practice that these physical tags include a written identifier and a machine readable, optical (e.g., QR code) or electromagnetic (e.g., RFID, NFC), element to facilitate reading. Any maintenance, repair, upgrade, trading, reuse and final decommissioning usually is associated with updates to an inventory system. Authorised users can read and modify the inventory, and each device can have an informative page associated to its digital identifier.



**Figure 1: the lifecycle of digital devices and data model**

In terms of the W3C Decentralized Identifiers (DID) architecture [W3CDID2021], see Figure 2, the digital identifier code for a device is a DID. These DID are currently referred as an OBIT in the context of Obada [OBIT]. A URL offered by the inventory service for each device (*DIDURL*) can allow to retrieve a specific informative page per device *DID DOC* as linked data (human readable HTML or other data format).



**Figure 2: W3C DID abstract architecture [W3CDID2021]**

When devices are decommissioned or *transferred* to a new organisation, that implies marking a device as not yet available or transferring their corresponding data. That includes a new user or even an intermediary organisation such as a refurbisher or a recycler too.

However, externals who do not own a device will not probably have access to any information about specific devices and their usage. Ownership of the physical device or having control over it implies access and control on its supporting data.

In a circular model, devices after the first and further use cycles can be collected by an IT asset disposition (ITAD) organisation. The ITAD collects the devices and performs a triage for either recycling or reuse. The refurbishment process consisting of an inventory, data erasure and perhaps an upgrade of the device. In eReuse (a project affiliated to Obada), this is done by the Workbench tool, that reports to a DeviceHub inventory server the serials collected among a group of devices (as in a pallet, called a “delivery note” in eReuse). Obada involves several stakeholders of the second-hand market of digital devices, many of them ITADs or related to that industry.

The list about relevant events about a collection of devices, including details (e.g., device identifier, operation, result, timestamp, agent, etc.) defines the history of these devices.

## 4.2 Context: safety properties

Different actors may be interested in this information for different purposes. While sometimes what matters is the information stored, other times what matters is certainty about the action done (proof, attestation). Like in accounting, a ledger is a permanent summary of all amounts entered in supporting journals which list individual transactions by date. This information can be of interest not only inside of an organisation but also by anyone willing to perform an audit or verify if and when any specific transaction took place.

The *global record of devices* (GRD) [Franquesa2015] is a verifiable record of accounting transactions about relevant events

Some of the main events about a device are the manufacturing, purchase, repair, upgrade or modification, decommission, transfer, data wipe, sale, recycling, loss. Reporting these events in an inventory can be useful to its owner, but in a circular economy, a device may go through multiple actors and organisations along its lifespan, and these may not trust each other, or may even risk colluding. Therefore, beyond the inventory systems for device owners in each organisation, a common verifiable data registry is needed to be able to record transactions and claims that affect any device along its complete lifespan.

The *verifiability* requirement translates into *irreversibility* of recordings (that operations already recorded cannot be undone or modified, sometimes referred to as *immutable* in the sense of *append-only*), and leads to data *replication* with ledger updates coordinated by a consensus majority decision, as a way to prevent any attempt of manipulation of ledger books.

However, irreversibility and the ability to be accessed by multiple actors, raises a requirement to preserve personal privacy and business confidentiality. Nothing in the ledger can be private or confidential, as it could not be removed without destroying a ledger completely. For that reason, most of the details about transactions in a ledger should contain verification information (e.g. proofs, hashes, signatures, timestamps) that enable the holder of any data to prove it was present or produced by the time a transaction was recorded: If a transaction record is a tuple (actor, signature, timestamp), *actor* can prove she had that data at the timestamp instant, as the data was hashed or signed by that actor, and that can be repeated now for verification by the data holder. The original data, the details, can be stored in any private inventory database (for device owners), while the ledger only stores the summary information about transactions for verification purpose. This global ledger log allows to search for events linked to a device along its lifespan across multiple organisations. That allows to verify traceability and impact information about devices. Combined with additional details of the data stored outside the ledger accounting books (such as private inventory databases), the data stored there allows us to generate verified circularity and social impact reports and metrics.

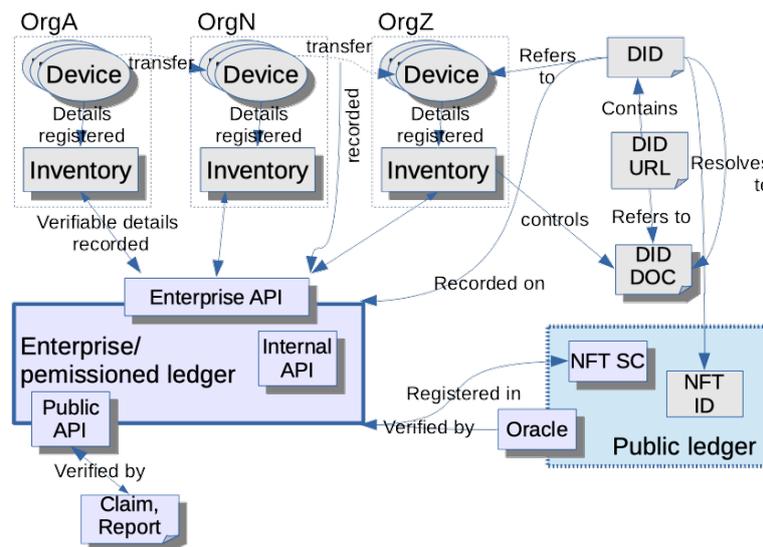
We may want to record commitments, such as the agreement that each operation requested has a cost (fee) to be paid to maintain the infrastructure, or that an economic deposit is left as a guarantee to return a device after a period of use, or to be given to a recycler by the manufacturer of that device (extended producer responsibility) by the time it is recycled. These commitments can be expressed as contracts, and ideally a ledger could account for details and produce side-effects like *inexorably* releasing funds whenever conditions are met. This idea corresponds to what a smart contract in Ethereum terms can do.

Inventories, as decentralised private registries that keep track of detailed digital information about owned digital devices, gets complemented with the verifiable registry or transaction ledger, accessible to a set of trusted actors (notaries, that record “annotations” for safety/verifiability), implemented as a distributed ledger. In general terms, there are three types of data about a device or a pool of devices or parts:



1. Private details (for owners, characteristics in an inventory),
2. Semi-public proofs about data safety (attestation, verifiability, inexorability in an internal/permissioned/federated accounting ledger) for notary-like roles,
3. Value or participation claims as a public record to attract economic investment to a pool of devices and data (public record of existence of devices, ownership, economic investment and charges/fees in a public accounting ledger).

The above is illustrated by Figure 3.



**Figure 3: Devices and their data: inventory, permissioned accounting ledger, public ledger**

The *eReuse testbed* is the *enterprise/permissioned ledger* (2) that works and interfaces with device inventory systems (1) as clients (Obada and API compliant device inventory management systems), and provides a permissioned (private to a federation of actors) accounting ledger for safety claims (proofs) about devices and service fees. It can interface (through a public API for checks, or an Oracle agent) to a public ledger (3) such as the public Ethereum mainnet (public cloud) where “real” money could flow linked to investment, management, trade and impact assessment of physical digital devices.

### 4.3 Data of interest in common

Data of interest refers to entities to be referenced both internally to the DLT (on-chain, for verifiability) and externally to the DLT but perhaps private (off-chain, for inventory details), and each side has their own formats. The internal format depends on the specific technology used for the verifiable registry (ledger), the external format depends on business domain requirements, expected properties and software used.

Identifiers related to digital devices and related data can be externally pointed to (*off-chain*) by unique identifiers (DID) to a specific inventory database, that can be referenced as URI/URL. In the DLT these will be represented by *addresses* of devices in a NFT contract



[ERC721][Obada2021] or a specific device smart contract address from a device factory, actors/wallets in an ERC20 contract, or *transaction IDs* of a recorded proof.

A device management platform/system manages information about devices (existence, transfer, lifecycle), actions (recording proofs of actions performed, such as data wipe or recycling), and charges (on fees for platform services or deposits). These systems are called DID controllers in the W3C did-core specification [W3CDID2021].

- About a **digital device**: DID for a device. That requires CRUD (*create, read, update, “delete” as recycle*) operations (can be enhanced by search and lookup operations)
- About an **action**: proof registration of different types. That requires CR (*create, read*) operations.
- About **fees**: participants can subscribe to future fee charges from the platform for direct debits as a result of platform operations. A participant can authorise future charges of (Obada) platform fees. A platform agent can request the charge of a service fee to a particular participant. If a subscription was recorded, the service fee will be directly debited to the participant account (wallet ID) in the ERC20 token contract in our permissioned Ethereum-based implementation. That requires CR operations.

The DID scheme, for the time being, can be as follows:

```
obit-did = "did:obada:" obada-specific-idstring
obada-specific-idstring = [ obada-class ":" ] obada-address
obada-class = "obit" / "proof" / "fee"
obada-address = (Hex-encoded specific unique identifier)
```

Information managed in a device management system about a digital device can be seen externally as a DID document providing details (e.g. JSON document), filtered by authorisation rules, about a specific DID URL, that contains a DID (DID subject or digital twin in Obada). Similarly for proofs and fees.

### 4.3.1 Devices

Computer devices (OBIT) IDs can be recorded in the DLT [OBIT]<sup>1</sup>. It is common to concatenate “manufacturer, part number, and serial number.” Obada proposes:

```
serial_hash = sha256(serial_number);
asset_hash = sha256(manufacturer + part_number + serial_hash);
version=0000;
Obit-DID=version + asset_hash + checksum
```

---

<sup>1</sup> <https://www.obada.io/standard/obit-formula>



### 4.3.2 Proofs

Proofs with a few details (including an ID) can be recorded in the DLT. The details to record are specific to each proof. No further details for brevity.

### 4.3.3 Fees

Fees (charges) with a few details (including an ID) can be recorded in the DLT, equivalent to direct debit charges. Those imply transfers among ERC20 wallets in our testbed Solidity Ethereum implementation of the DLT (from the service customer to the service provider, to pay for platform services).

### 4.3.4 Documents

Document proofs can be registered in our DLT. A hash/signature of the document with a timestamp, not the content of the document as such, is recorded in the DLT. Later on, it can be verified: that the document was seen and a proof (hash/signature) was recorded at a given previous time.

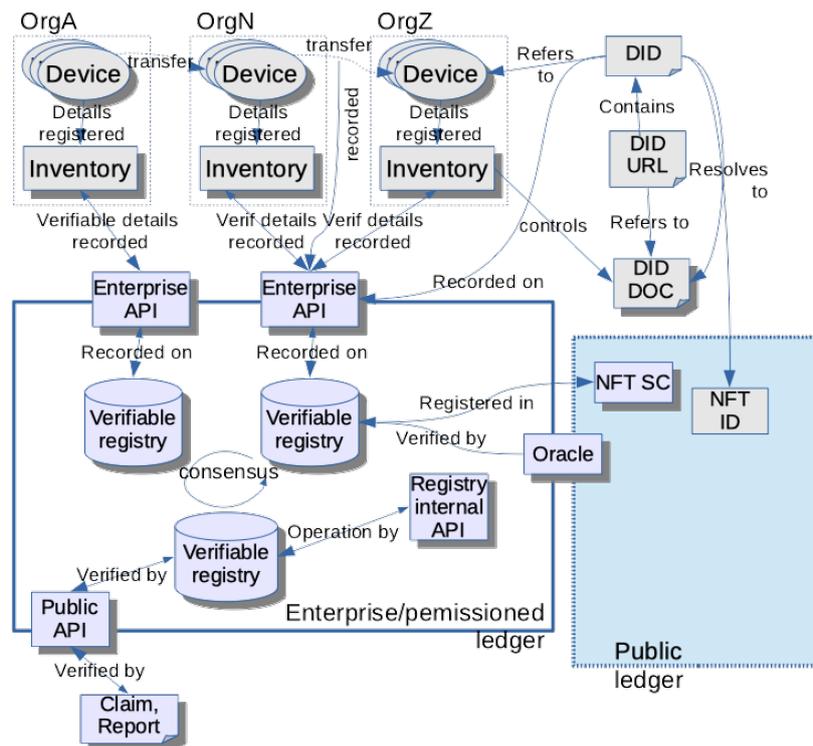
### 4.3.5 Reports

The DLT allows queries to verify registered items –devices, proofs, fees, even documents (hashes/signatures of them)–. We can generate reports (e.g. a PDF document) that can be shared with the public. These reports can be later on verified against the ledger entries over a public API.

## 4.4 Architecture

Our system architecture is composed by a device management platform that involve several federated service-elements or service instances (*Obada proof-of-concept prototype, the DeviceHub system integrated in the Obada architecture*) with a DLT/blockchain backend (permissioned ledger) offered by the eReuse testbed, and the possibility of a separate public ledger presence. The detailed system architecture is illustrated in Figure 4.





**Figure 4: System and data architecture**

The device management platform manages devices and offer/manage/comply-with DID namespaces that relate to (twinning) material digital devices and parts, and provide metadata (documents) and methods for those authorised. For each device, there is a **unique identifier** that can be referenced as a DID. That DID, an OBIT, can be recorded in a “Verifiable Data registry” through our DLT API testbed.

In OBADA, an employee of a partner organisation uploads a **spreadsheet** with info about computers (generated manually or somehow automated). A **Client helper** tool or **node** computes the OBIT IDs from the data about each device (serials) and invokes the enterprise **DLT API** to record each OBIT in the shared verifiable data registry. Client helper instances can provide as well an OBIT **resolution** service.

The DLT API testbed provides a *verifiable data registry* with a backend, that in our testbed implementation is composed by a DLT (Ethereum PoA using 3-5 *geth* instances, with *Solidity* smart contracts) and a database cache (Postgres) which has tables of data for quick lookup, as a cache updated by events on the blockchain. It keeps a record in the internal (DLT as addresses) about external references (DID). We use in the testbed a Prometheus daemon to collect logs and a Graphana instance for queries and visualization of log data about usage and the experiments.

**Open source software release:** the software has been released as a public repository of the research group: [https://gitlab.com/dsg-upc/ereuse\\_dlt\\_api](https://gitlab.com/dsg-upc/ereuse_dlt_api) that is related to other repositories. The code for the project spans several other repositories:

- [eReuse Smart Contracts](#): Core smart contracts of the project.
- [Token transfer system](#): An ERC20 token compliant smart contract system.
- [Ethereum node deployment](#): Deployment of a node of our testbed.
- [API tests](#): Scripts to test the performance and correctness of the testbed.

The testbed software is licensed with an Affero General Public License.<sup>2</sup>

## 4.4 Experimentation

We have worked with our US partners to explore a verifiable registry API that could be provided by our DLT backend and that could be the basis for a public specification (standard) for the traceability of devices and the accountability/verifiability of recorded transactions. We have adapted our testbed to provide API operations that can keep a record of verifiable information as the basis to comply with the requirements of the scenario of use. This is related to the alternative implementation of a QLDB API in the Obada prototype.

Furthermore, we have performed a cost (overhead), benefit (performance) and scalability (load profile and service limits) analysis and evaluation to assess the feasibility and scalability of a realistic system built along this design.

## 4.5 API

The aim of a RESTful API is to enable other actors and their systems to interact with the eReuse Ledger (DLT) experimental testbed. The data to be recorded in our distributed ledger relates to the lifecycle and verifiability of digital devices, proofs about important actions, and fees charged for platform services. This data can be identifiers as addresses or URI, and details as summaries (hashes) of documents.

The distributed ledger records these transactions reliably (irreversible, append only), and (inexorably) executes smart contracts to manage the catalogue of devices and associated economic deposits, proofs about device actions that can generate impact reports, and ERC20 wallets that store balances of units of value (tokens) and can be used to manage deposits (associated to devices as guarantee for commitments, such as returning or recycling a device), and fees (charged for platform services).

Expected actors, as clients, can be digital device management platforms like the Obada prototype client, including the USOdy DeviceHub application. The aim is to standardise the Ledger API and produce an open-source reference implementation to validate the API specification for standardisation organisations such as ISO or ITU-T.

---

<sup>2</sup> [https://en.wikipedia.org/wiki/GNU\\_Affero\\_General\\_Public\\_License](https://en.wikipedia.org/wiki/GNU_Affero_General_Public_License)



Our experimental API back-end is based on ERC specifications such as ERC20 [Ethereum2021] for a token system and the ERC721 [0xProject2021] equivalent for non-fungible items and deposits, the eReuse proof system [Franquesa2019][Franquesa2020], and for the front-end the W3C DID specifications [W3CDID2021] for a common naming across the on-chain and off-chain sides.

The implemented API of the testbed used for the experiments is as follows:

Endpoint	Parameters
/api/devices/:deviceAddress GET - Get a device's info.	- deviceAddress: (URL parameter) - Ethereum address.
/api/devices/create POST - Create a device.	- uid: (String) - initValue: (Number) - owner: (String) - <i>This owner represents the registrant (hash) instead of the owner Ethereum address. The Ethereum address is defined in the API.</i>
/api/devices/transfer POST - Transfer a device's ownership.	- deposit: (Number) - new_owner: (String) - Ethereum address. - new_registrant: (String) - device_address: (String) - Ethereum address.
/api/devices/generateProof POST - Generate a proof related to a device.	- proof_type : (String) - device_address: (String) - Ethereum address. (The rest of the parameters depend on the proof type.)
/api/devices/recycle POST - Recycle a device.	- device_address: (String) - Ethereum address.
/api/devices/createStamp POST - Create a stamp on a hash.	- hash: (String) - SHA3-256
/api/devices/checkStamp POST - Check if a stamp exists on a hash.	- hash: (String) - SHA3-256

Although we currently use “deviceAddress” as device identifiers in several operations to simplify the performance evaluation of the testbed, we could incorporate other identifiers as they consolidate, such as the OBIT identifiers used in the Obada client, or in general a decentralized identifier (DID) [W3CDID2021]. In fact, the CREATE operation establishes this relationship between an external representation (DID) and the internal as Ethereum addresses in our DLT implementation. These API endpoints are a superset and currently equivalent to the QLDB API implemented in the Obada prototype [OBADARD]. Syntax differences are due to implementation details of each code, can be addressed with glue code, and solving them is part of the development of final API specifications.

Regarding the ERC20 wallet compatible operations, these are accessible through an internal smart contract interface instead of the previous REST API, as these are limited to storing deposits and the payment of service fees for verifiability and accountability services.

## 4.6 Testing

First of all, we have performed tests to verify the correct operation of the resulting system. After simple unit tests that verify the correctness of each API operation, we have performed a test that exercises a typical sequence of operations involved in a realistic situation.

Testing has been based on client-side scripts that interact with the DLT or API and server-side metrics collection of every node using a Prometheus/Grafana setup. We aimed to provide validation of the various functions that can be performed by the system and to check its performance under different scenarios.

To validate the correctness of the operations provided by the API, we have designed a test that performs calls in a logical order that could represent a real world situation. The sequence of calls is as follows:

1. Create a device instance with the owner being the Ethereum account that the API manages.
2. Generate a proof on the device.
3. Transfer the ownership of the device to another Ethereum account.
4. Try to transfer the same device to another account after having lost the ownership. This is expected to fail.
5. Try to recycle the device after having lost the ownership. This is also expected to fail.
6. Recycle an owned device (creates a new owned device and tries to recycle it).
7. Create a stamp on a SHA3 hash.
8. Check the existence of said stamp after creating it.
9. Try to create a stamp on the same hash expecting to fail.
10. Try to look for the existence of a stamp on an unstamped hash, also expected to fail.

By the end of the development, this test could execute properly, meaning that the experimental API and the collection of smart contracts were working as expected.

## 4.7 Experimental performance analysis

To check for performance under different conditions we have gathered data across **three** different experiments:

1. Performance of a *geth* node under different transaction load.
2. Performance of the API node (HTTP or websocket) with an increasing operations load.
3. Rate of transactions to collapse a *geth* node.



For our **first** experiment, we wanted to stress the system with batches of increasing number of transactions. Our goal here was to measure how the blockchain performed with an increasing transaction load. For this purpose, we conducted the experiment several times for each batch size, and calculated the median of all the sampled values. We decided to use the median rather than the average after observing the variability of the samples. This variability exists due to the unpredictable spikes in network latency that come from running the tests remotely through the internet.

This test was carried out by Adrian Manco as part of his master's thesis [Manco2021], where he evaluates the implementation of an ERC20 token trading system for the testbed. The transactions were sent directly to one of the Ethereum nodes (without passing through an external API) and involved simple trading operations between different accounts.

We started performing these transactions between the same two accounts and quickly found some anomalous behaviour in the system. Provided that the transactions sent to a single *geth* node always come from the same Ethereum address, and given a high enough number of transactions (around 300), the node seems to sometimes fail to propagate every single one. Due to the on demand (0 block time) proof of authority consensus algorithm of our nodes, and the fact that the transactions from a single account have to be executed in a specific order (indicated by a nonce value). We found that sometimes the node that had to close the next block did not receive the transaction that had to be executed next so it was not able to do it. This left the blockchain in a locked state that would not execute the transactions nor generate new blocks unless we interacted with it again. To go around this issue, we ended up performing the transactions across a much larger number of accounts which effectively solved the problem.

One thing to take notice of, is the possibility of sending these transactions through HTTP or a WebSocket interface. We found that if we used HTTP, a lot of requests started to get dropped after a burst of around 200 transactions. Seeing that this did not happen through WebSocket, we decided to use this protocol during this experiment.

Finally, we decided on an upper bound of 1000 transactions per batch. This was because we started observing anomalous behaviour in the system beyond this point.



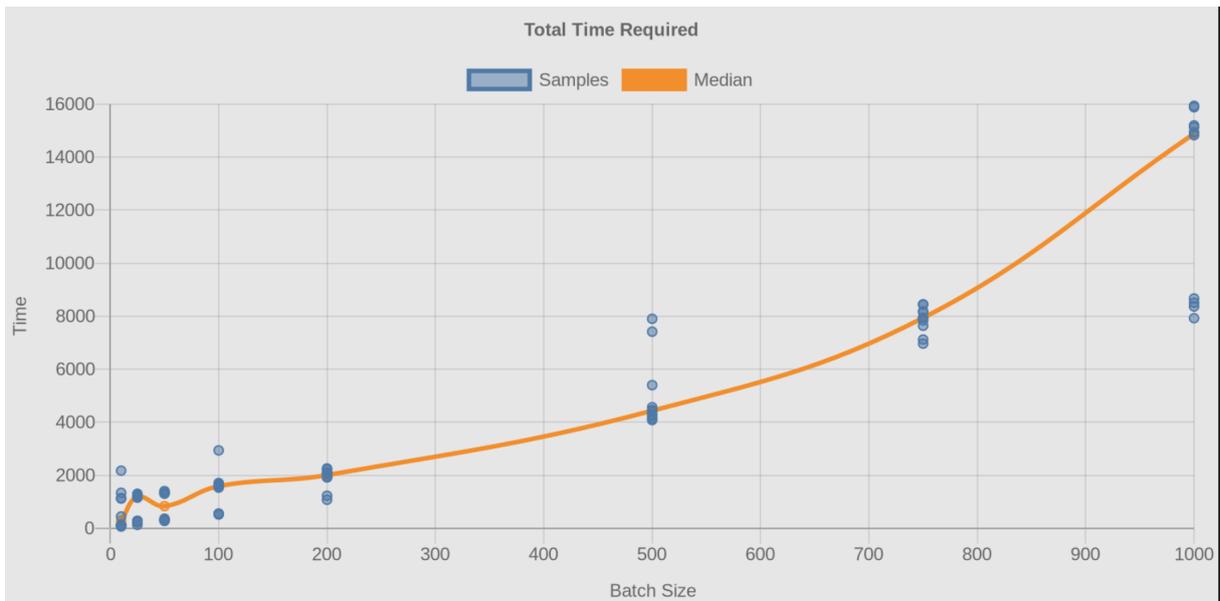


Figure 5: Time required for each batch

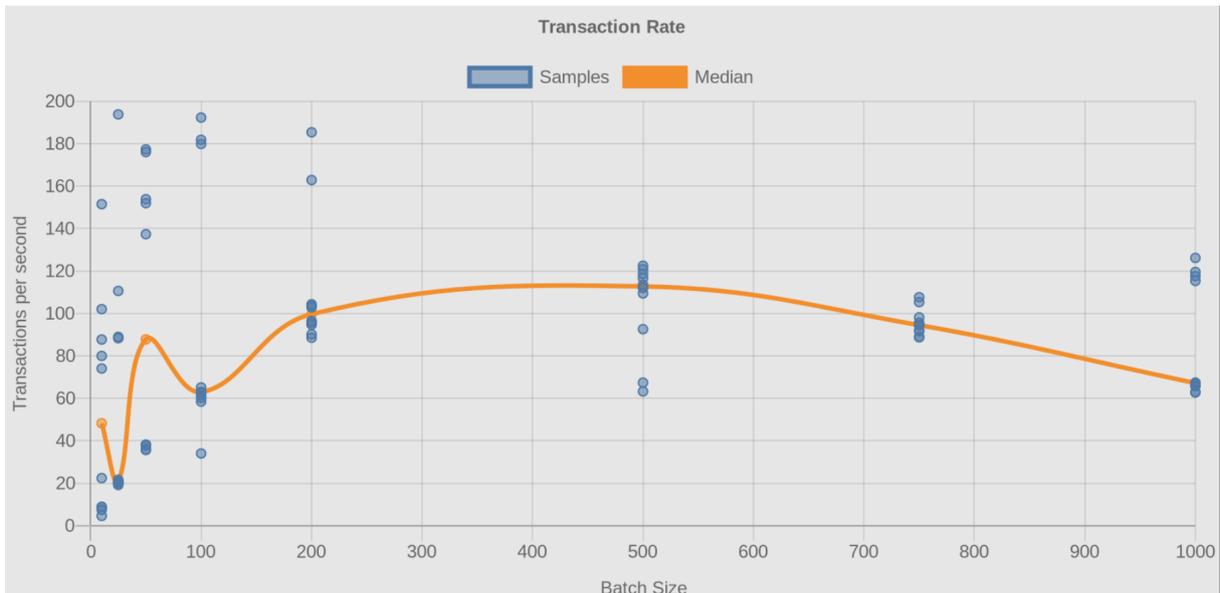


Figure 6: Rate of transactions per second



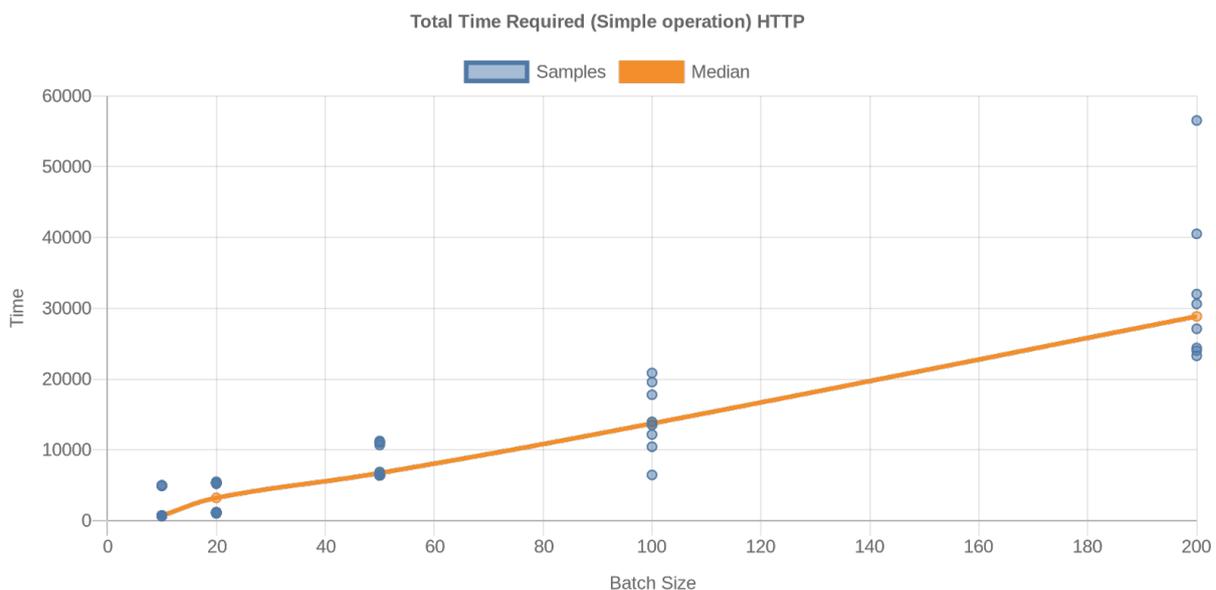
**Figure 7: Resource usage in the first experiment**

As we can see in the charts, the X axis represents the size of each batch of token transactions, while the Y axis represents the specific metric measured in each chart. The sampled measures can be seen coloured in blue on the charts. After calculating the median of the observed samples, we obtained an aggregated Y value that we included in the graphic, and linked them so the trend was easier to understand. Regarding the resource usage, we took advantage from the Grafana dashboard that the team built, and that was fed with the data provided by the Prometheus monitoring system.

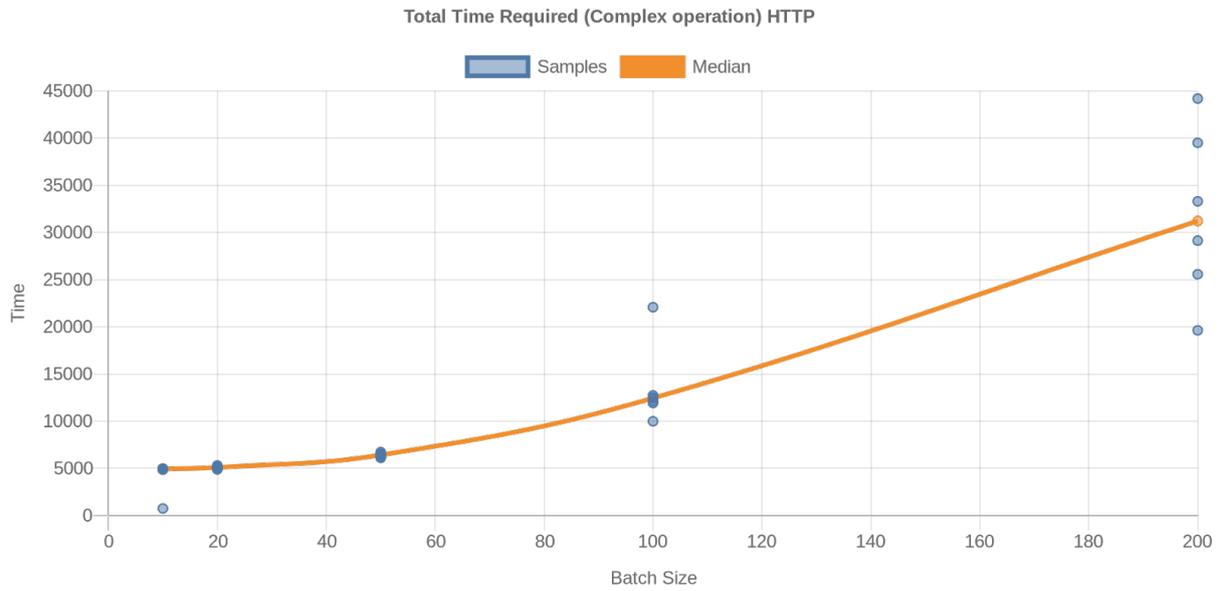
For our **second** experiment, we wanted to perform a similar test but this time passing through the external API. This poses some problems as the API, by design, can only use a single Ethereum account. This fact meant that we would not be able to avoid the propagation problem found earlier, and limited us to an upper bound of 500 transactions per batch.

The calls to the API are always made through HTTP as we determined this would be the most convenient protocol for this use case. Nevertheless, the API can still communicate with the blockchain through HTTP or WebSocket, so we chose to perform the test using both protocols, as to determine which would be the most convenient.

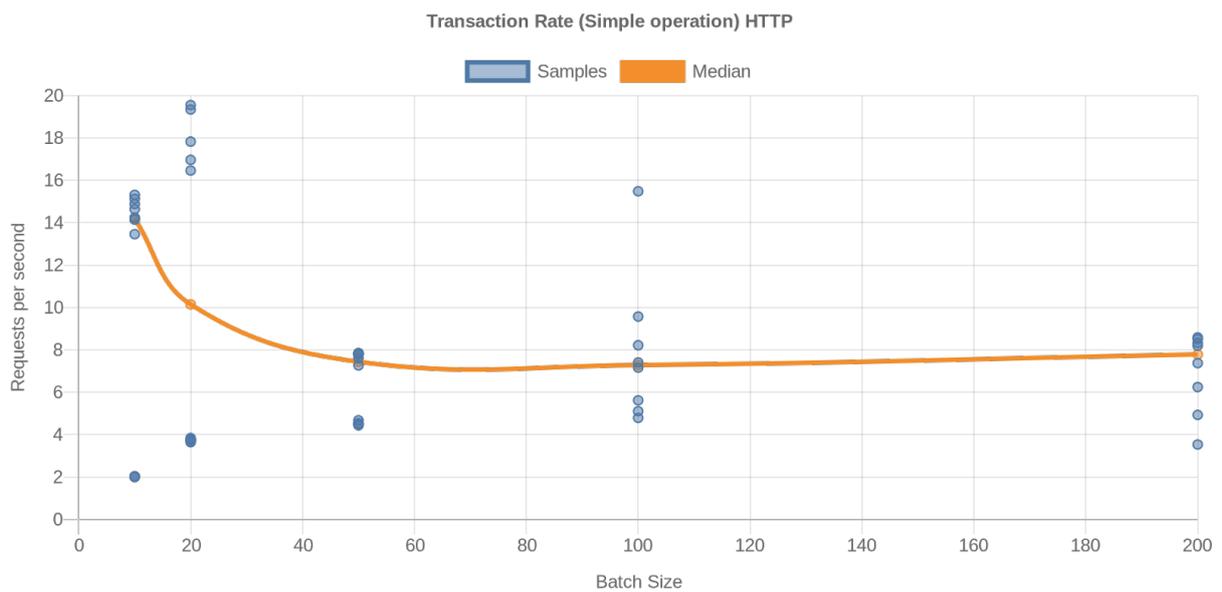
Because this test was going to execute operations related to a more complex smart contract, the difference on computing complexity of the different functions was actually really noticeable. We deemed necessary and interesting to execute the test with the most and least complex ones, as to see the difference in the response of the system.



**Figure 8: Total time for simple transactions using HTTP.**

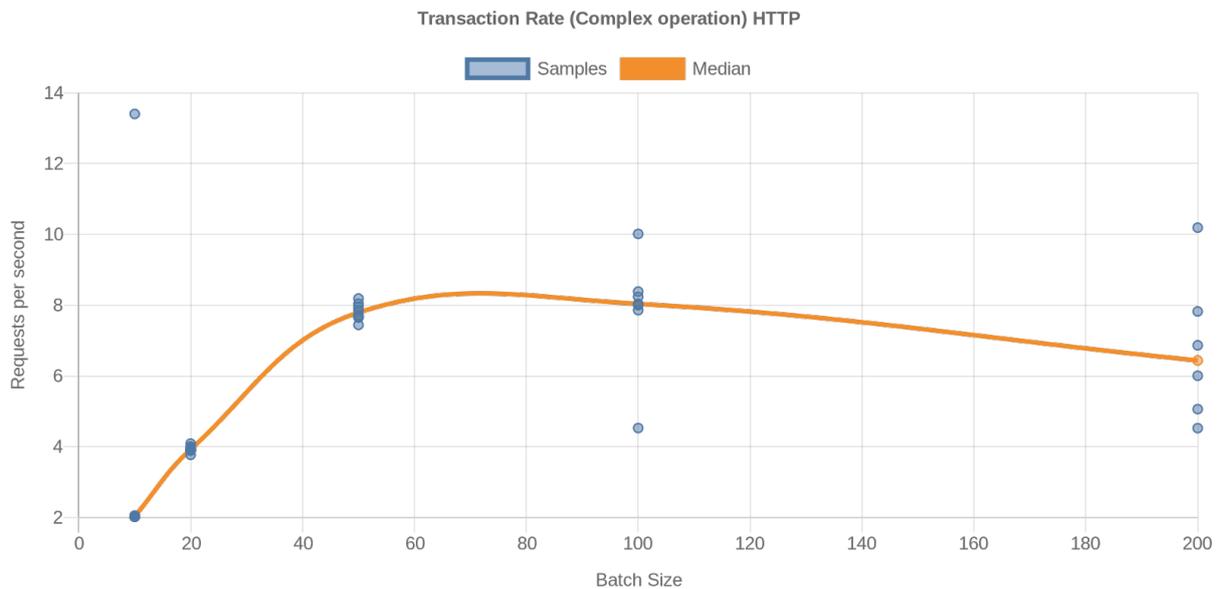


**Figure 9: Total time for complex transactions using HTTP.**



**Figure 10: Transaction rate for simple transactions using HTTP.**





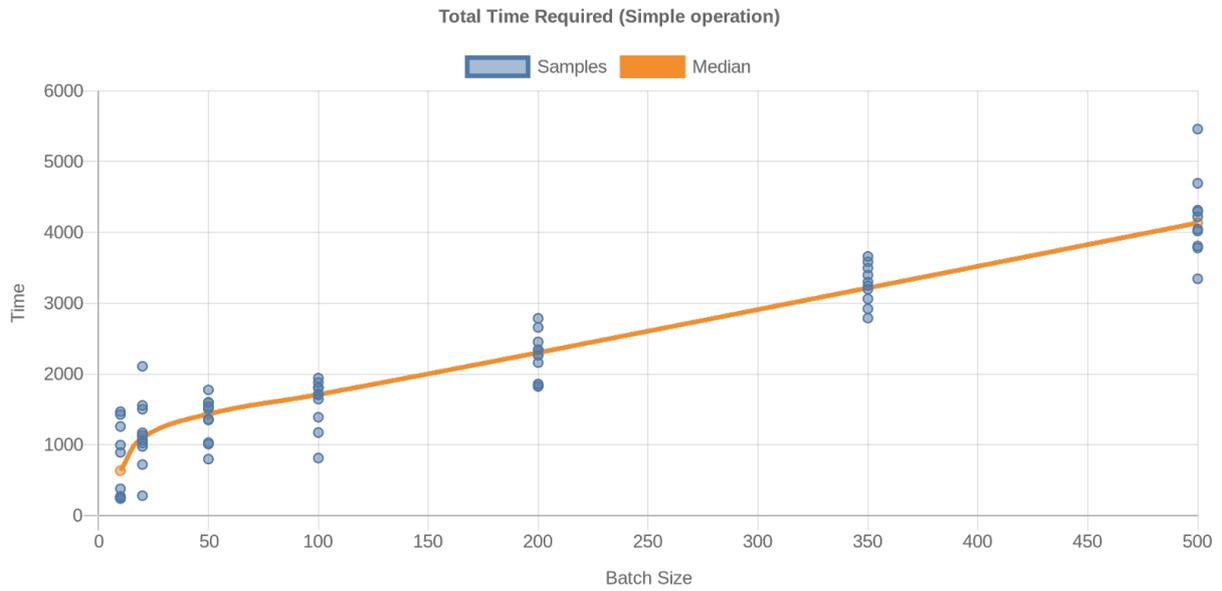
**Figure 11: Transaction rate for complex transactions using HTTP.**

These are the results over the HTTP protocol. The upper bound is reduced to 200 over this protocol, as we encounter the problem described in the first experiment, of the Ethereum node dropping requests beyond this point.

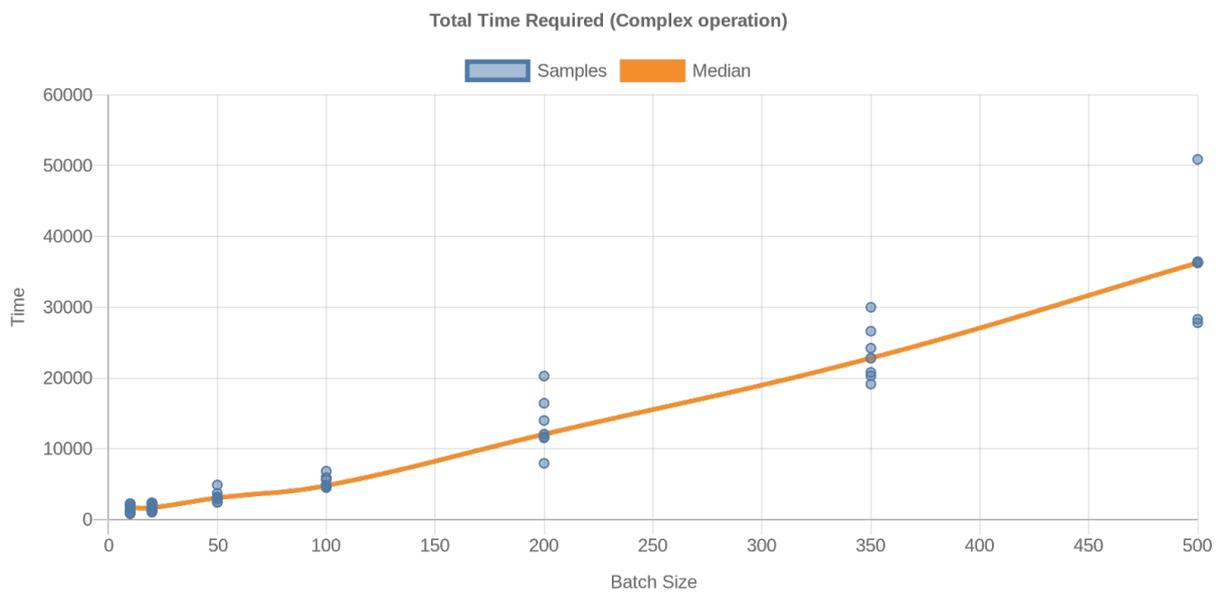
The simple transaction is a call that recycles a device. This transaction costs around 30000 gas units. Gas represents the amount of computational power required to execute the transaction on an Ethereum system. Furthermore, this transaction costs around the same amount of gas as the transactions used in the first experiment, which would allow us to directly compare the results.

The complex transaction is a call that creates a new device. This consists, among other things, of creating a new smart contract that represents this device. The gas required for this operation is around 5 million, making it the most complex of them all, requiring around 167 times the gas used on the simple transaction.



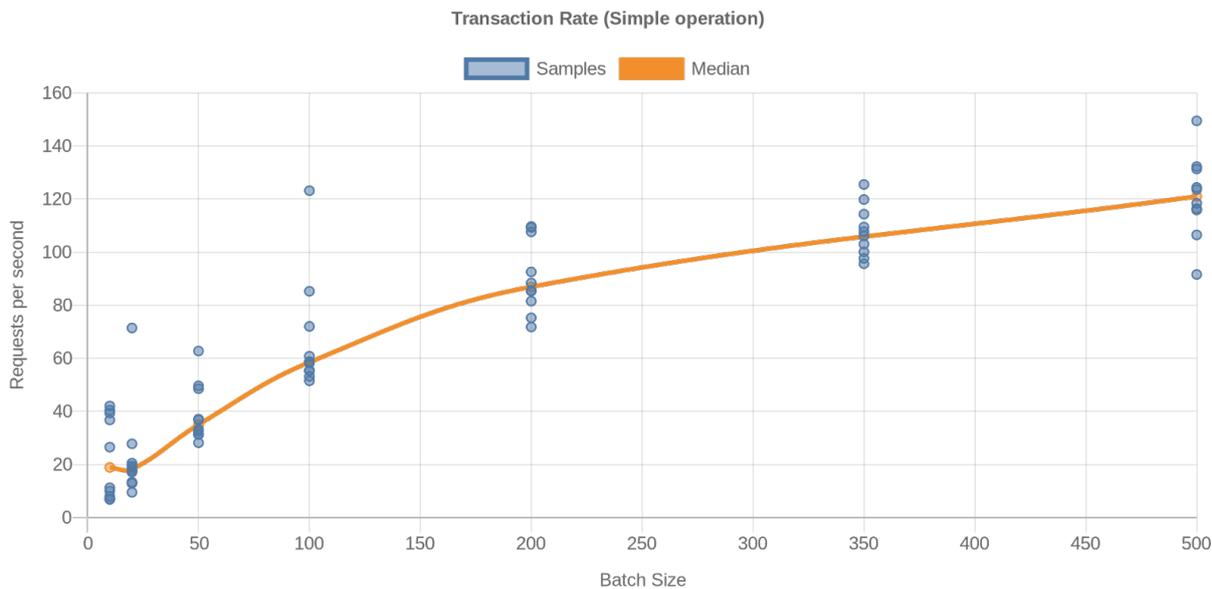


**Figure 12: Total time for simple transactions using WebSocket.**

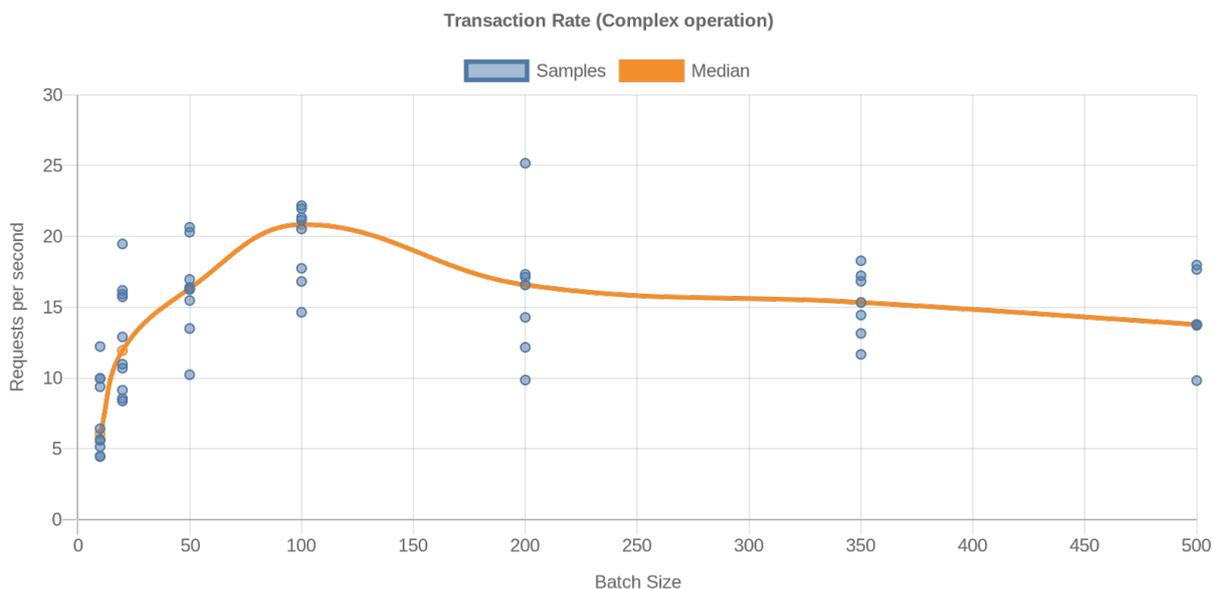


**Figure 13: Total time for complex transactions using WebSocket.**





**Figure 14: Transaction rate for simple transactions using WebSocket.**



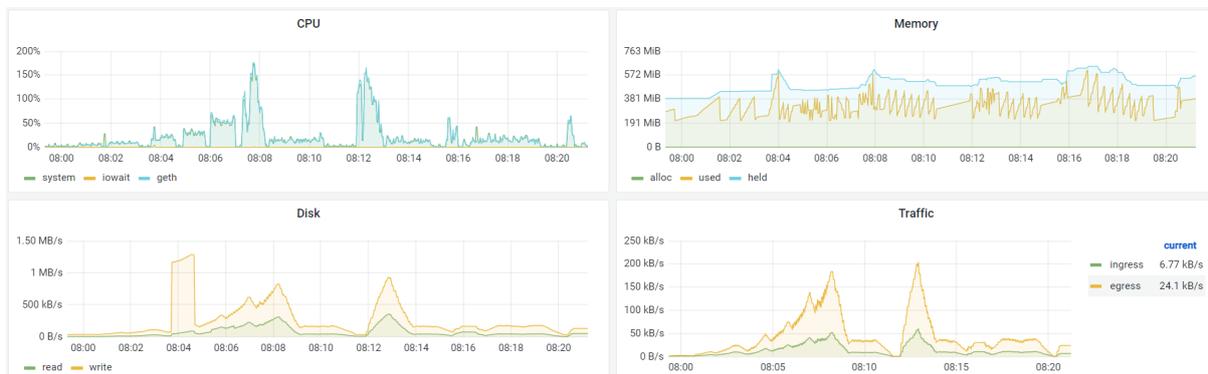
**Figure 15: Transaction rate for complex transactions using WebSocket.**

These are the results for the WebSocket protocol. Every result has noticeably improved when compared to the HTTP version.

For our **third** experiment, we wanted to assess what was the maximum number of simple transactions that the blockchain could process per second **without a server collapse**. With this in mind, we decided to send bursts of transactions, separating the beginning of said bursts by one second, but not controlling when they ended. In effect, what this produced was a steady generation of bursts, and as the size of the burst grew larger, the blockchain started overlapping one set of transactions with the next. The idea was testing at which point the processing power of the blockchain became insufficient. We would know that by seeing in the



server’s log either a high resource usage, or by seeing anomalous behaviours as a consequence of the collapse of the equipment. We first used HTTP as transport protocol with results in Figure 12.



**Figure 16: Resource usage in the third experiment (HTTP)**

And afterwards repeated the same exact experiment using WebSocket transport instead with results in Figure 13.



**Figure 17: Resource usage in the third experiment (WebSocket)**

## 5. Discussion and Analysis on Results

With the previous validation tests, we can prove that we achieved the objectives set at the start. A device can exist, be identified and owned in our system. Events regarding these devices can be properly recorded as proved by our proof generation and hash stamp functions. Devices can also contain any kind of information stored as a hash inside their metadata, to allow to **safely identify** any third party clients that access the system through one of the entities that controls an instance of the API.

As for performance, we will go through each experiment separately.

Regarding the **first** experiment, in a real scenario similar to the one we assessed, the system's performance would be high enough to meet the needs. The time function suggested that the



evolution of the curve was super-linear, with an increasing slope as the size of the batches was incremented. Furthermore, the rate function was somehow symmetrical, although it featured some interesting differences. For smaller batch sizes, it showed high degree of variance, probably due to the way the blockchain closed the blocks. As the size of the bursts became larger, the values began stabilising, and we could see that the rate of transactions per second decreased more with each increment in the burst size. The resource usage provided by Grafana suggest that the load did increase in every possible aspect during the execution of the experiments, but was far from reaching its maximum value. The maximum CPU usage measured was of the 142% out of an available 400%, since we had 4 cores available in the node, and the memory, disk and traffic also showed a remarkably good performance, far from compromising the availability and timeliness of the system. Finally, it is worth mentioning that at the end of our experiment, the values measured suggested a quick decrease in the performance for bigger batch sizes. This leaves the door open to considering how large a batch of transactions the system would be able to process without collapsing.

Regarding the **second** experiment, we can clearly see the superiority in performance of the WebSocket protocol over HTTP. In both time and transaction rate we can observe that the results are sometimes up to 10 times better. On top of that, the upper bound of around 200 transactions in the case of HTTP, could mean that the system may not be adequate for certain situations. It may be a good idea for the API to communicate to the blockchain via WebSocket, which would be invisible to the final user as they would still make calls to the API via HTTP. Nevertheless, this could vary by case, as HTTP is a stateless protocol that does not leave a connection channel open at all times, which could be an advantage if the number of transactions sent to the blockchain is not too high.

As for the comparison of complex and simple transactions, we again observe a high **variance** on the smaller batch sizes. However, as the values stabilize, we see how the system gets better results on the simpler transactions, although it is important to notice how these results do not scale proportionally with the difference in gas between the complex and simple transactions. This illustrates how much overhead the whole process adds in relation to the capabilities of the EVM. It is highly noticeable in the HTTP version of the experiment, where both kinds of transactions get very similar results, indicating that the overhead introduced by HTTP is even higher. This could be due to the fact that a connection needs to be opened for every single transaction.

This experiment also wanted to assess how the external API in the middle of the client and the blockchain would affect the performance. We quickly realized that this would be very dependent on the machine this instance would run on, so we decided to perform the calls in the same machine that was running the API. This resulted in the API not being a problem at all and giving us a very similar result to the first experiment, seeing as those transactions, and the simple transactions used in the second experiment were almost equivalent in gas.



Regarding the **third** experiment, the resource usage observed in the Grafana dashboard confirms a maximum effective throughput of approximately 200 simple transactions per second using HTTP. This gives enough room for a successful system behaviour, but should be taken into account when studying the production scenario, the system would need to face. If the average expected load is supposed to be close to that value, there is the risk of service outage, so further measures have to be taken with regards to this. If the load is expected to be lower, the risk is somehow relative, but having such a defined limitation is a symptom that the system still needs to be polished and optimised, or the use case revised. The previous is derived from the readings provided by Grafana, as well as the fact that the own dashboard started malfunctioning when we arrived to a certain threshold. In the resource usage plots we can see how the CPU reaches a maximum at around the 150% of the capacity. That corresponds to the experiments with a burst size of 200 transactions. The processing power used had been growing steadily (as can be seen by the different valley areas in the plot), but after 200 transactions per second, the *geth* process collapsed, the logging of the testing script started showing errors, and the entire dashboard crashed. If we go back to the results from the first and second experiments, we will see that some outliers indicated a maximum of around 200 transactions per second, which is consistent with the situation being described here. From these facts and results we can conclude that there is a breaking point in the performance of the system, and that its correct workflow cannot be guaranteed with such a load.

With WebSocket, Grafana showed a substantially different behaviour. The resource usage grew in a similar way than it did with HTTP, but after bursts of 200 transactions, the *geth* process did not collapse, nor did Grafana. We were able to keep observing the charts evolving, and even increase the burst size to 750 and 1000 transactions without any critical failure appearing. While significantly positive, the values also showed that, after 200 transactions per burst, the CPU consumption trend reversed, the system began consuming less CPU overall as we increased the independent variable, and the logging showed a clear overlap between bursts. That might be caused by the sequentiality with which the *geth* process itself processes the incoming requests, or a need to revise the architecture and configuration of the distributed Ethereum ledger.

For the third experiment, we could safely conclude that the current state of the infrastructure limited the correct functionality to 200 transactions per second. Beyond this threshold, the system began showing undesirable behaviours with both HTTP and WebSocket.

Future work may include the performance-cost evaluation of the effect of higher latencies and number of *geth* instances, as well as different DLT node characteristics.

## 6. Present and Foreseen TRL

Regarding TRL of API development and interoperability of the DLT backend. We have gone from 4 (tech validated in lab) to TRL 5 (tech validated in relevant environment) through the



experiments, and TRL 6 (tech demonstrated in relevant environment) through the participation of multiple OBADA actors in the experiments (U. of Nevada-Reno, Tradeloop, USODY organisations, OBADA developers). Some aspects could reach TRL 7 (system prototype demonstration in operational environment) but we have not focused on security, availability and access control aspects to allow an open demonstration in an operational environment.

## 7. Exploitation, Dissemination and Communication Status

**Exploitation and dissemination** as input to standardisation processes for APIs in circular economy forums related to DLT, such as ISO TC 307 or ITU-T Q7/SG5:

Regarding the slow progress with **ISO**, we have discussed the Obada new global ISO work item proposal with the Spanish standardization organisation (UNE) that has agreed to vote in favour. We are waiting for the outcome of the global voting for it.

**ITU-T** approved the L.GDSPP "Requirements for a global digital sustainable product passport to achieve a circular economy" work item<sup>3</sup> [GDSPP] in the SG5 virtual meeting 11-20 May 2021, with Leandro Navarro as editor with the additional support and willingness to contribute from Orange, Huawei and Cisco. This work items results from the work of this project.

A Liaison statement was sent from ITU-T to the European Commission (DG CNECT, DG ENV, DG GROW).

The first internal draft of the recommendation was presented and discussed in an online meeting of L.GDSPP in July 8, with positive feedback from the participants from Huawei, UN Basel convention, and Apple. The document is being developed with further progress meetings expected in September and following months. The start of a new work item and contributions to an ITU-T recommendation is considered a major achievement for this project.

**Dissemination and exploitation as open source software:** the software has been released in the public repository of the research group: [https://gitlab.com/dsg-upc/ereuse\\_dlt\\_api](https://gitlab.com/dsg-upc/ereuse_dlt_api). As with previous software releases, the testbed software is licensed with an Affero General Public License.<sup>4</sup>

---

<sup>3</sup> Public news by ITU-T about the work item:

<https://www.itu.int/en/myitu/News/2021/07/15/16/04/New-ITU-standards-project-to-define-a-sustainability-passport-for-digital-products>

<sup>4</sup> [https://en.wikipedia.org/wiki/GNU\\_Affero\\_General\\_Public\\_License](https://en.wikipedia.org/wiki/GNU_Affero_General_Public_License)



**Further research** as a funded project in the Trublo 1<sup>st</sup> open call (8/2021-3/2022): We aim to investigate trust and reputation models built on permissioned blockchains to improve the transparency, accountability and provenance of user supplied data about the lifespan of digital devices to report multimedia portfolios of documents that confirm transactions and reward personal or organisational climate change efforts.

**Non-profit stewardship organisation** to fund, maintain, consolidate and evolve the software and specifications of the testbed beyond the reported experiments: UPC is part of the **eReuse.org** initiative, a federation of organizations [Franquesa2016] working on digitized management of digital devices in a circular economy model, as well as part of the Obada Foundation [ObadaF2020] and the OBS software company.

We plan to create **USOdy**, a spin-off software and services company of the research group at UPC, with graduated members of the current research group and collaborators, as a member organization of the Obada Foundation and the OBS software company, that can exploit the DeviceHub application integrated with DLT services based on the testbed software.

**Further research and development at UPC** in the framework of the eReuse.org initiative with reference code and a pilot with Obada and eReuse stakeholders, as well as contribution to standardisation.

#### **Dissemination:**

Communications as selected publications planned in diverse stakeholder forums: academic, industry, policy, social.

An **academic publication** is planned after the NGI Atlantic project based on the experimental results, and parts of this report.

Obada (Rohi Sukhia and Ronald Lemke, our USA partners) plans to **present** in the ITAD Summit 2021<sup>5</sup> conference the outcomes and demo of the current demonstrator and user-interface prototype, a verifiable ledger for second-hand digital devices, that involves among other components the NGI Atlantic testbed.

We co-organized and held a **session on circular economy in the Eurodig** conference, on Tuesday 29 June 10:30 CEST as part of the Greening Internet Governance Part II [EuroDig2021]. The session was called “*Circular and digital: Internet governance as part of the solution*”, facilitated by Leandro Navarro and Beat Estermann from the University of Applied Sciences of Bern.

---

<sup>5</sup> ITAD Summit 2021 - Huntington Beach August 18-19 – Conference, <https://www.itadsummit.com>



We contributed a section about the digital product passport in the Association for Progressive Communications (APC.org) guide for the circular economy for activist organisations. The guide will be launched in September – October 2021.

### Communication:

Presentation to the OBADA Board of Directors of the progresses, as well as achieved and expected outcomes in June 15, 2021. A summary:

- The NGI Atlantic testbed: status of the DLTs to keep track of devices DID/Obits, the API, proofs and certificates.
- How NGI Atlantic relates to the Obada-tech discussion and the integration of all for a demo in architectural terms.
- Which features each module provides, and how that functionality we can either shown working or show tech feasibility.
- How Obada can contribute to the work in ITU-T L.GDSPP standard in SG5 and ongoing discussion with SG20.
- Clarify the relation between UPC, USOdy as spin-off company of UPC, and the open licensing/contribution from our software and specs to OBS/Obada.

## 8. Impacts

**Impact 1:** Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation.

Policy cooperation enhanced through the work on the global UN ITU-T L.GDSPP initiative with additional interest and support from Orange and Huawei European organizations, as well as Apple or CISCO from US. Support through the UNE Spanish standardisation organization in the voting for ISO to start a new related work item. Collaboration and expanded participation in the Obada initiative with US partners but also interest from additional EU partner organizations. These recommendations/standards can help policy concertation across both regions.

This policy cooperation is particularly relevant in the context of climate change, the IPCC recommendations, the ITU-T L.1470 translation to the required reductions in the ICT sector, that translate into policy initiatives like the Digital product passport (in Europe lead by the EC, and globally lead by the ITU-T L.GDSPP work item started in this project), as well as the European Green Deal and equivalent Green New Deal initiatives around the world<sup>6</sup> at the same time there are initiatives for post-pandemic recovery that build on green innovation.

---

<sup>6</sup> [https://en.wikipedia.org/wiki/Green\\_New\\_Deal](https://en.wikipedia.org/wiki/Green_New_Deal)



**Impact 2:** Reinforced collaboration and increased synergies between the Next Generation Internet and the Tomorrow's Internet programmes.

Our US research partner at UNR works in life-cycle sustainability for computing particularly how it pertains to topics in computer systems research, but NSF considers that interest is not compatible with the referenced DCL. They are exploring “sustainability” programs (sustainable computing) at NSF to enhance the US research side. No results yet as the program has not yet developed completely. However, the exploration with NSF has raised awareness about the topic there, and we expect resulting in enabling future collaboration with our US partners.

**Impact 3:** Developing interoperable solutions and joint demonstrators, contributions to standards.

Continued and closer collaboration beyond the limits of the NGI Atlantic project, through coordinated experiments that lead to refinement of the business model, integrated demonstrator systems, and the development of APIs for interoperability.

Definitely this work leads to interoperable solutions and standards (linked to the L.GDSPP work item started, potentially and eventually an ISO work item if approved). We had discussions about joint API specifications and interoperability tests to improve the performance and functionality of the current prototype demonstrator.

**Impact 4:** An EU - US ecosystem of top researchers, hi-tech start-ups/SMEs and Internet-related communities collaborating on the evolution of the Internet.

The EU-US ecosystem has expanded through the participants in the eReuse (more than 20 social organizations), Obada<sup>7</sup>/OBS and ITU-T L.GDSPP participants.

## 9. Conclusion and Future Work

The NGI eReuse-Ledger testbed is a permissioned distributed ledger to support experimentation about device traceability and related verifiability aspects. The NGI Atlantic support has enhanced the EU-US collaboration, as it has enabled the development of interoperable specifications, solutions and joint demonstrators with complementary contributions from both regions. It has contributed to start and feed a new work item for standardisation in ITU-T and potentially in ISO.

---

<sup>7</sup> <https://www.obada.io/participants> member organizations: 5 trade associations, one certification body, 2 authorities, 2 blockchain companies, 17 application software providers, 6 user organizations, 10 ecosystem stewards.



Future work consists on expanding the functionality of the testbed and its API (with support from the NGI Trublo project), contribute to have a system demonstrator with expanded functionality that can motivate different stakeholders in industry and society for digitized accountability of digital devices in the context of environmental sustainability and climate change mitigation, optimization of the testbed components for better performance and scalability, and finally contribution of the lessons learned to public specifications as with the work in ITU-T or ISO.

In summary, this project has contributed to expand and consolidate a network of top researchers, hi-tech start-ups/SMEs and Internet-related communities collaborating on the evolution towards a more social, inclusive and environmentally friendly Internet, for the benefit of people and the planet, a survival challenge for our society.

## 10. References

- [OxProject2021] "Non-Fungible or Unique Tokens on the Ethereum Blockchain." OxProject.  
<http://erc721.org>
- [ERC20] ERC-20 Token Standard, (Visited URL 2021),  
<https://ethereum.org/en/developers/docs/standards/tokens/erc-20/>
- [ERC721] ERC 721 - OpenZeppelin Docs, (Visited URL 2021),  
<https://docs.openzeppelin.com/contracts/3.x/api/token/erc721>
- [Ethereum2021] "ERC-20 TOKEN STANDARD." Ethereum.org.  
<https://ethereum.org/en/developers/docs/standards/tokens/erc-20/>
- [EuroDig2021] Greening Internet Governance, Part II – Enabling an Environmentally Sustainable Digital Transformation in Europe – FS 01 2021,  
[https://eurodigwiki.org/wiki/Greening\\_Internet\\_Governance,\\_Part\\_II\\_-\\_Enabling\\_an\\_Environmentally\\_Sustainable\\_Digital\\_Transformation\\_in\\_Europe\\_-\\_FS\\_01\\_2021#PART\\_I%3A\\_Input](https://eurodigwiki.org/wiki/Greening_Internet_Governance,_Part_II_-_Enabling_an_Environmentally_Sustainable_Digital_Transformation_in_Europe_-_FS_01_2021#PART_I%3A_Input)
- [Franquesa2015] D. Franquesa, et-al. Breaking barriers on reuse of digital devices ensuring final recycling. In EnviroInfo and ICT for Sustainability 2015. Atlantis Press, 2015 (Best paper award) <https://www.atlantis-press.com/proceedings/ict4s-env-15/25836176>
- [Franquesa2016] David Franquesa, Leandro Navarro, and Xavier Bustamante. 2016. A circular commons for digital devices: tools and services in ereuse.org. Second Workshop on Computing within Limits (LIMITS'16) ACM.  
<https://dsg.ac.upc.edu/node/914>
- [Franquesa2019] Franquesa, D., and L. Navarro. 2019. "An IT Asset Disposition Platform That Incentivises and Certifies the Circular Economy of Digital Devices: Operative MVP - December 2019." UPC.  
<https://dsg.ac.upc.edu/sites/default/files/dsg/deliverable-ereuse-mvp.pdf>.
- [GDSPP] Proposed new work item on "Requirements for a global digital sustainable product passport to achieve a circular economy", <https://www.itu.int/md/T17->



- [SG05-210511-TD-GEN-1828](https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=17130), Work item entry: [https://www.itu.int/ITU-T/workprog/wp\\_item.aspx?isn=17130](https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=17130)
- [ITU-T-SG] ITU-T Study groups (Study Period 2017-2020), <https://www.itu.int/en/ITU-T/studygroups/2017-2020/Pages/default.aspx>
- [Küsters2010] Ralf Küsters, Tomasz Truderung, and Andreas Vogt. 2010. Accountability: definition and relationship to verifiability. In Proceedings of the 17th ACM conference on Computer and communications security (CCS '10). Association for Computing Machinery, New York, NY, USA, 526–535. DOI: <https://doi.org/10.1145/1866307.1866366>
- [Ledger2020] “An IT Asset Disposition Platform That Incentivises and Certifies the Circular Economy of Digital Devices: MVP Testing – February 2020.” UPC. <https://dsg.ac.upc.edu/sites/default/files/dsg/ledger-del2-feb.pdf>
- [Manco2021] A. Manco Sanchez and Navarro, L., “Blockchain-based system for subscription payments in circular economy model”, 2021. Master Thesis, UPC. <https://dsg.ac.upc.edu/node/931>
- [ObadaF2020] Obada Corporate Documents, 2020, <https://www.obada.io/docs/>
- [ObadaNFT2021] The Decentralized OBIT NFT Registry for Physical Assets, 2021, <https://www.obada.io/token-proposal/obit-nft-registry.html>
- [OBADARD] OBADA Reference Design: An inventory manager application, (Visited URL 2021), <https://dev.rd.obada.io>
- [OBIT] Obit Formula, (Visited URL 2021), <https://www.obada.io/standard/1-1-1>.
- [q7/5rep] Draft Report of Question 7/5 (Virtual, 11-20 May 2021) <https://www.itu.int/md/T17-SG05-210511-TD-GEN-1733>
- [W3C2019] “W3c JSON-LD 1.1: A JSON-Based Serialization for Linked Data.” W3C. <https://w3c.github.io/json-ld-syntax/>
- [W3CDID2021] W3C. Decentralized identifiers (dids) v1.0: Core architecture, data model, and representations, 2021. <https://www.w3.org/TR/did-core/>
- [W3CUC2021] W3C. Decentralized identifiers (dids) v1.0: use cases, 2021. <https://www.w3.org/TR/did-use-cases/>

## 11. Glossary

API	Application programming interface
DID	Decentralized identifier
DLT	Distributed Ledger Technology
eReuse	Electronics Reuse
EVM	Ethereum Virtual Machine
GDSPP	Global digital sustainable product passport
ISO	International Organization for Standardization
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
NGI	Next Generation Internet



OBADA	Open Blockchain for Asset Disposition Architecture
PoA	Proof of Authority
R&D	Research and Development
TRL	Technology Readiness Level
UPC	Universitat Politècnica de Catalunya
W3C	World Wide Web Consortium
WIT	Waterford Institute of Technology (Coordinating Partner)

