Review Open Call SME Stage 1 Experiments

GoldenOwl

Daniele Miorandi

U-Hopper srl

Review meeting

All over the world, 11/11/2020

A 100s intro

GoldenOwl in a nutshell

https://www.youtube.com/watch?v=oRFDBSquhac&feature=youtu.be&ab_channel=U-HopperSrl

Experiment description

Background and motivation

- UH is developing a new digital product, based on DLTs, for the management of education & training certificates
- Based on the use of a permissioned blockchain for storing digitally signed copies of certificates
    - Tamper-resistant
    - Robustness & resilience (thanks to decentralisation)
    - Users in control of their own certificates
- Prototype available (used for a PoC with customer in late 2019). Yet doubts on its scalability (performance and economics).

Concept and objectives

- Goal: assess experimentally the scalability of our solution
- Technical objectives:
    - To design an experiment for assessing the scalability of our solution using Grid'5000
    - To package the existing GoldenOwl software stack to easily and quickly deploy instances and execute tests in an automated way
    - To perform scalability tests of GoldenOwl on Grid'5000 and to collect the relevant experimental data;
    - To analyse the experimental data and understand bottlenecks

Experiment setup

- Experimental facility used: Grid'5000
    - Nancy gros cluster: 124 nodes (1 CPU Intel Xeon Gold 5220, 18 cores/CPU, 96GB RAM, 447GB SSD, 894GB SSD, 2 x 25Gb Ethernet)
- Experiment automation: python script + enoslib + ansible
- Resources used:
    - One control node (scripts+InfluxDB)
    - One orderer node (DLT, Telegraf)
    - Two web server nodes (WS/APIs, Telegraf)
    - From 2 to 50 peer nodes (DLT, CouchDB, Telegraf)

Experiment setup

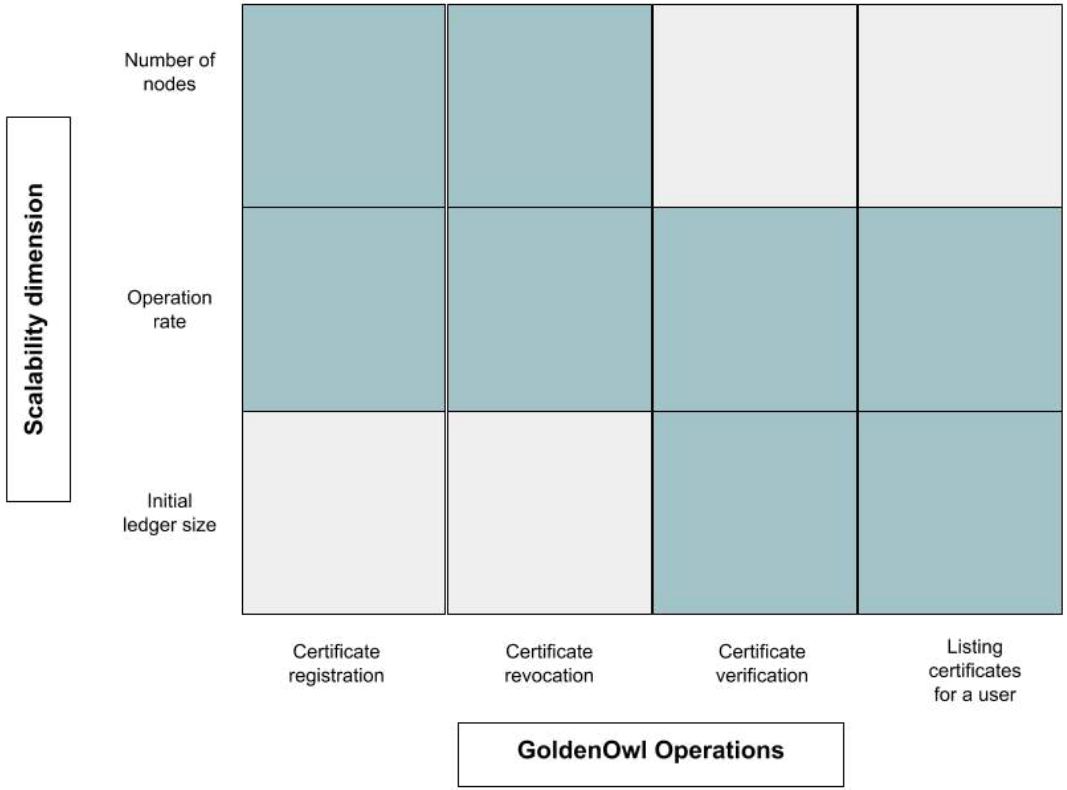

Project results

Measurements

	Experiment-1	Experiment-2	Experiment-3	Experiment-4	Experiment-5	Experiment-6	Experiment-7	Experiment-8
Operation	Certificate Registration	Certificate Registration	Certificate Revocation	Certificate Revocation	Certificate Verification	Certificate Verification	Listing Certificates	Listing Certificates
Number of nodes	2, 10, 20, 50	10	2, 10, 20, 50	10	2	2	2	2
Operation rate per second	100	1, 10, 100, 1000	100	1, 10, 100, 1000	1, 10, 100, 1000	100	1, 10, 100, 1000	100
Initial ledger size	100	100	10000	10000	100	100, 1000, 10000	100	100, 1000, 1000

- 10 runs per experiment

Measurements (cont'd)

- 223 individual runs, collected around 8.08 GB of raw data.
- Some examples (see final report for full records):

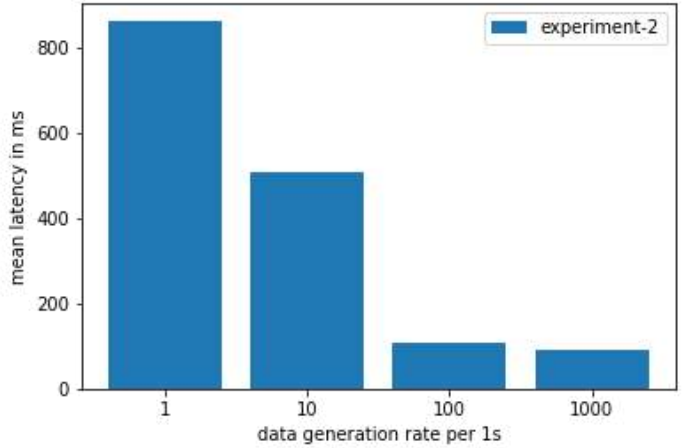

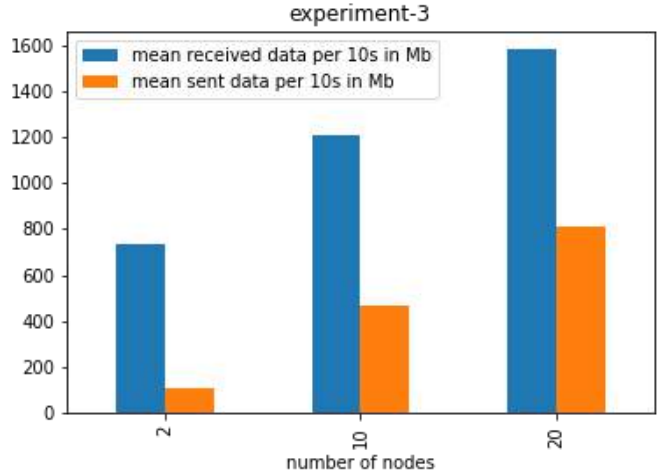

Lessons learned

- Latency decreases when the operation rate increases – counterintuitive, due to setting of batch timeout and batch size parameters
- Latency increases when the number of node increases – superlinear, due to consensus protocol execution time
- The initial ledger size does not influence the latency
- CPU usage is not influenced by the number of nodes but is influenced by operation per second rate
- The software is not ready to scale to a really high number of nodes – with 50 peer nodes experiment does not complete – could be influenced by single orderer
- Network traffic increase with number of nodes – potential bottleneck in terms of unit economics
- Memory usage is stable

Business impact

Value perceived

Main value: identification of the scalability bottlenecks of our current implementation

Other values perceived:

- ability to run experiments in a distributed, large-scale setting
- knowledge of a set of tools for easing/automating deployment and data collection (in particular Enoslib)

Value perceived (tech perspective)

- Objective data on performance
    - At a scale we would not have been able to do by ourselves
- Benchmarking & reproducibility: be able to single out the impact of one parameter on system scalability
    - Showing the unexpected impacts of some system configuration parameters and implementation choices

Value perceived (business perspective)

- Lower risks related to premature go-to-market
- Speed up identification of bottlenecks and issues for the tech team
    - (Modulo time to acquire knowledge of how to work with the experimental facility & automate experiments)

What's next?

- Tackle the identified scalability issues → (Technical) product development roadmap
    - 4-5 months of time
    - 10-12PMs of effort estimated
- Do it in an experimentally-driven, agile fashion: quick build/deploy/measure loops
- Will keep on doing it on Fed4FIRE+: GoldenOwl2.0

Feedback

Used resources and tools

- One single experimental facility used: Grid'5000
    - Did not leverage federation aspects
- Tools: Enoslib + Ansible

Added value of Fed4FIRE

- Fed4FIRE+ provided us with:
    - Grid'5000, a cloud-in-vitro: realistic yet controlled environment (important for benchmarking
    - Access to large-scale infrastructure: not economically feasible otherwise for the company
    - Access to knowledge: documentation + role of patron
    - Set of tools for (partially) automating experiments

Issues encountered

- Scheduling of large number of nodes on Grid'5000 (up to 54 nodes simultaneously) was problematic
    - Maintenance (not always planned)
    - High-priority jobs taking over

Suggestions for future improvements

Our focus: full experiment automation

- Tools: Enoslib is a very powerful tool for experimenters - current support within Fed4FIRE+ federated facilities is limited?
- Methodology/knowledge: tutorials, examples and best practises on how to automate experiments lifecycle

Q&A

This project has received funding from the European Union's Horizon 2020 research and innovation programme, which is co-funded by the European Commission and the Swiss State Secretariat for Education, Research and Innovation, under grant agreement No 732638.

WWW.FED4FIRE.EU