Review Open Call SME Stage 2 Experiments

GoldenOwl2.0

Daniele Miorandi

U-Hopper srl

FEC9

All over the world, 27/05/2021

Experiment description

Background and motivation

- **GoldenOwl**: a new product for the management of education & training certificates
- Based on the use of a permissioned blockchain for storing digitally signed copies of certificates
    - Tamper-resistance
    - Robustness & resilience (thanks to decentralisation)
    - Users in control of their own certificates
- Prototype available, yet doubts on its scalability → Stage-1 experiment (Jul-Sep 2020)
- Stage-1 experiment led to the identification of scalability bottlenecks → definition of a new product development roadmap

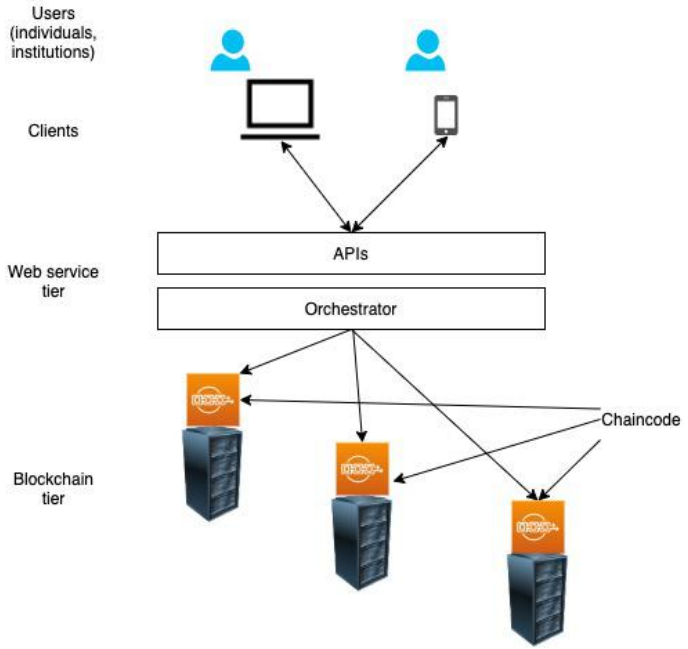

Concept and objectives

- **Goal**: to improve the scalability of GoldenOwl
    - Follow experimentally-driven approach, based on short build/deploy/test loops (sprints).
- **Technical objectives**:
    - To define a series of build/deploy/measure sprints, each one characterised by (i) a scalability issue (ii) one (or multiple) change(s) in the GoldenOwl software (iii) a detailed test specification
    - To execute the sprints, each one resulting in a new software release, in the execution of a set of tests on Fed4FIRE+ and in the analysis of experimental data
    - To package the final result into a new version of GoldenOwl
    - To provide Fed4FIRE+ with actionable feedback on issues/limitations as well as benefits obtained

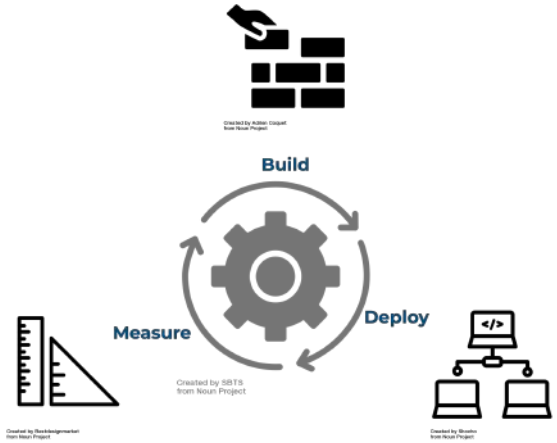

Experiment setup

- Experimental facility used: **Grid'5000**
- Resources used:
    - Nancy Gros cluster, each machine has an Intel Xeon Gold 5220 CPU and 96 GB of Ram, we used up to 57 machines.
    - Rennes Paravance cluster, each machine has 2 x Intel Xeon E5-2630 v3 CPU and 128 GB of Ram, we used up to 54 machines.
- Experiment automation: custom Python script + EnOSlib + Ansible

Experiment setup

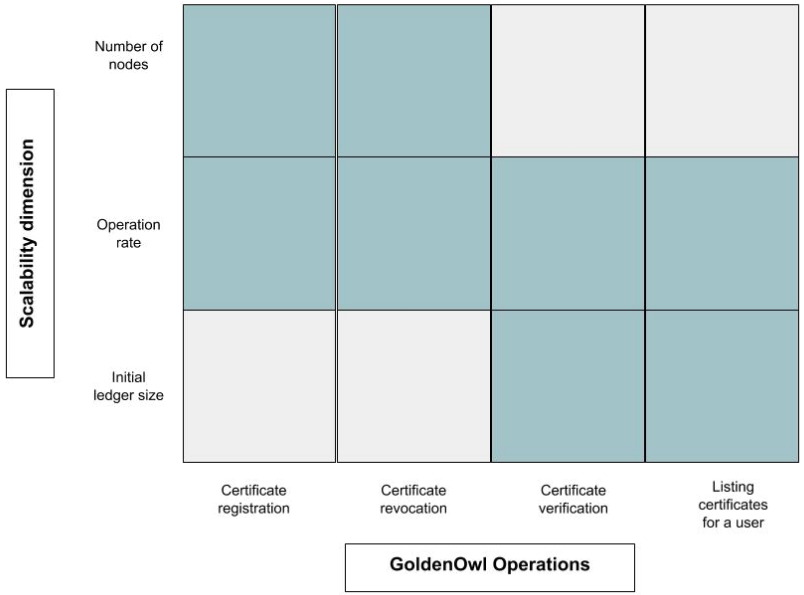

Test-1 Operation involved: certificate registration. Number of nodes: 2, 10, 20, 50. Operation request rate: 100. Initial ledger size: 100.
Test-2 Operation involved: certificate registration. Number of nodes: 10. Operation request rate: 1, 10, 100, 1000. Initial ledger size: 100
Test-3 Operation involved: certificate revocation. Number of nodes: 2, 10, 20, 50. Operation request rate: 100. Initial ledger size: 10,000.
Test-4 Operation involved: certificate revocation. Number of nodes: 10. Operation request rate: 1, 10, 100, 1,000. Initial ledger size: 10,000.
Test-5 Operation involved: certificate verification. Number of nodes: 2. Operation request rate: 1, 10, 100, 1,000. Initial ledger size: 100.
Test-6 Operation involved: certificate verification. Number of nodes: 2. Operation request rate: 100. Initial ledger size: 100, 1000, 10,000.
Test-7 Operation involved: listing certificates for a user. Number of nodes: 2. Operation request rate: 1, 10, 100, 1000. Initial ledger size: 100.
Test-8 Operation involved: listing certificates for a user. Number of nodes: 2. Operation request rate: 100. Initial ledger size: 100, 1000, 10000.

Project results

Measurements

- Initially: re-done all Stage-1 experiment with updated GoldenOwl software, new hack for optimising experiment duration and 50 nodes
- Four sprints
    - Batch timeout/batch size
    - Fan-out
    - Endorsement policy
    - Number of orderers
- For each sprint, different experiments, 10 runs per experiment
- 20+ GB of experimental data collected and shared on Zenodo

Measurements (cont'd)

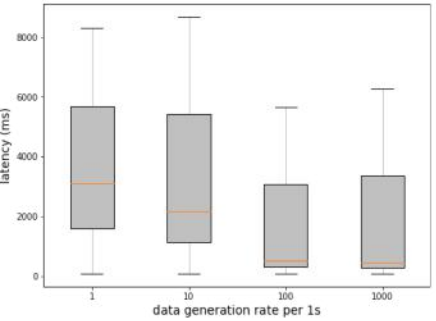

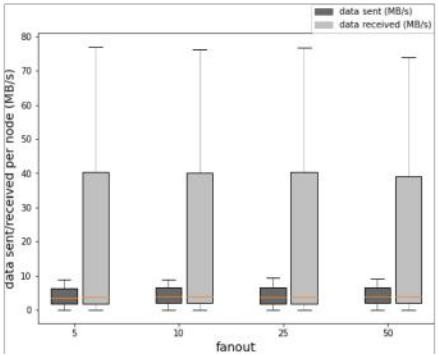

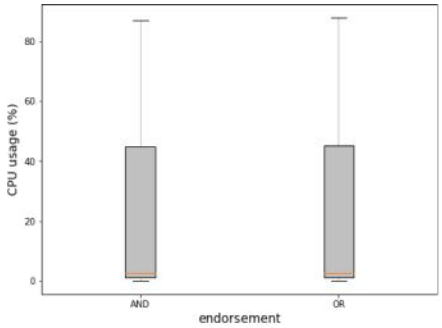

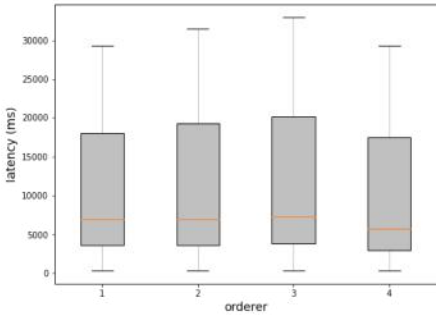

Lessons learned

- Smaller batch timeout reduces latency; smaller batch size reduces variability
- Reducing the fan-out reduces traffic and CPU while not impacting latency
- Changing endorsement policy reduces traffic while slightly increasing latency
- Increasing the number of orderers reduces latency

- *Latency reduced by 15% (up to 60% in low-requests scenarios)*
- *Traffic reduced by 20%*
- *CPU usage reduced by 15%*

Business impact

Value perceived

Two main values:

- *Speed up* the product development process (est. 2.5x)
- *Derisking* related to scalability issues once in production

Value perceived (tech perspective)

- Access to a *cloud in vitro*
    - Objective data on performance within repeatable settings
    - Benchmarking and the ability to isolate the impact of single changes in the software codebase
    - Supporting a *scientific* approach to software enhancement

Value perceived (exp. facilities)

Own servers Small scale, repeatable experiments	**Fed4FIRE+** Large scale, repeatable experiments
Small scale, non-repeatable experiments	**Public clouds** Large scale, non-repetabale experiments

Value perceived (business perspective)

- Save time (and related costs) in the development of a digital product
- Lower risks related to premature go-to-market
    - Ability to test at scale in a controlled setting

What's next?

- GoldenOwl considered TRL-5
- Current focus: business modelling & lean validation
- Next: commercial PoCs
    - Test value proposition
    - Test software framework in operational conditions
- We would like to keep on using Grid'5000 ('digital twin') after the end of the project

Feedback

Used resources and tools

- One single experimental facility used: Grid'5000
    - Did not leverage federation aspects
- Tools: EnOSlib + Ansible (+ custom Python scripts)

What worked well

- Managed to run experiments with large number of nodes
    - Was problematic in Stage-1
    - Thanks to feedback/insight by patron (INRIA)
    - No more issues with high-priority jobs taking over
    - (Required a hack due to lack of support for container jobs in EnOSlib)
- Manage to complete all experiments on time
- Excellent results in terms of performance enhancements of GoldenOwl

What could have worked better

- We run long tests, and we needed to run *a lot* of tests
  - Iterative nature of our workplan → stop-and-wait
  - Availability was in some cases problematic (no servers available, run out of quotas)
  - Long electrical maintenance in Nancy in Feb. → Move to Rennes (slowdown)
  - Erratic behaviour in Apr. in Nancy

Long-term feedack to F4Fp

- Need consistent SLAs if you plan to move to a commercial offering
- Build a network of tech partners expert in infrastructure for running Experiments-as-a-Service (EaaS)
    - Potential opportunity to build spinoffs by F4Fp partners

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