BELA

GOALS

In the Bela proposal, our primary objective is to conduct multiple performance experiments under realistic conditions. The data produced will allow us to configure the platform and optimize the overall performance as well as the footprint over the underlying infrastructure optimally for less costly operations.

CHALLENGES

Systems that support operations of such scale, are required to perform

- under uncertain conditions of failure,
- while maintaining data redundancy support.

DEMO SETUP

Virtual Wall

- 1 x Node with 8GB RAM and 8vCPU for Master-node
- 15 x Nodes with 4GB RAM and 4vCPU for our Cluster (Worker-nodes)
- 4 x Nodes with 4GB RAM and 5vCPU for our End-customers
- 1 x Public IPv4 assigned on the Master-node

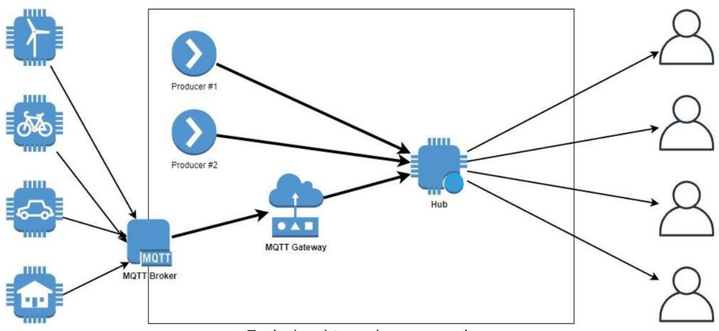

Typical end-to-end user scenario

City of Things

- 5 x nodes with 4GB ram and 4vCPU for the virtual IoT devices
- All nodes include a public IPv4 by default allowing them to be accessed directly

RESULTS

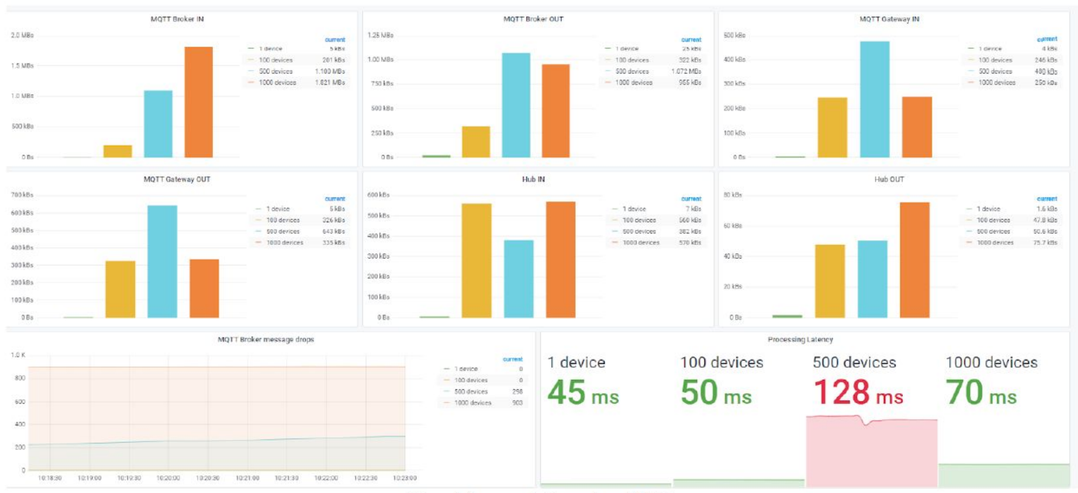

Dashboard for the KPIs

We used Grafana to visualize the gathered KPIs.
Selected KPIs are:

- Processing Latency
- Broker Throughput IN
- Broker Throughput OUT
- Gateway IN
- Gateway OUT
- Hub IN
- Message drops

MORE RESULTS

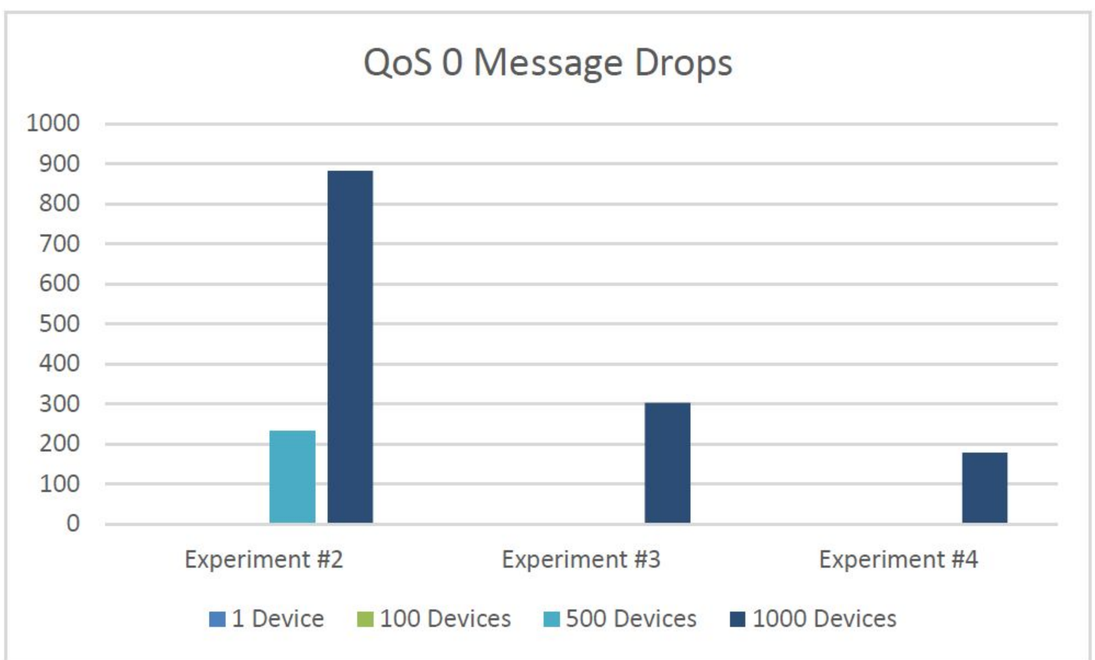

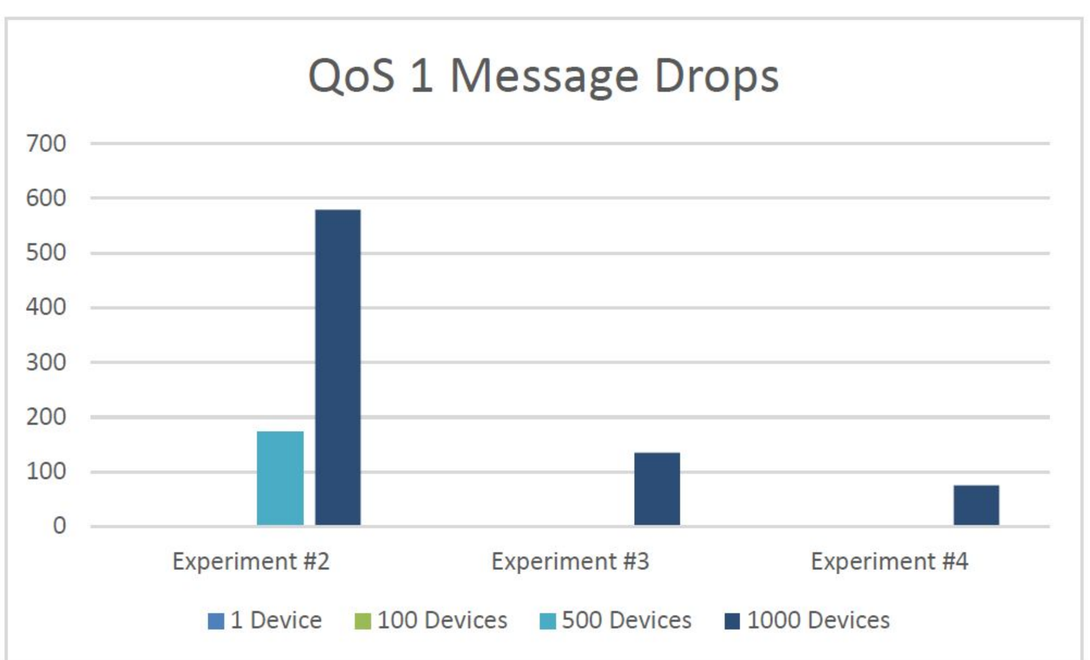

CONCLUSIONS

- The first set of experiments was a remarkable learning experience for Anadyme members, where we have managed to:
  - Improve platforms stability
  - Enriched Lavva.io with multiple new features
  - Pivoted our business model to a pay-per-use model
  - Managed to successfully stress-test our platform
  - Through extensive product discovery we have managed to create a clear roadmap plan
  - Improved platform performance and mission critical KPIs

POST MORTEM

We need to do further experimentation of the following issues:

- Stress test Lavva.io platform under chaotic conditions,
- Add support for new cluster types such as MQTT/RabbitMQ
- Implement Federation support with multi-clustering capabilities
- Implement auto-scaling as well as auto-healing capabilities
- Scale-out and optimize the capabilities of the platform on more VMachine types