

GOALS

In order to be able to scale our platform we need to assure that one core part of the architecture is reliable and scalable for our purposes, specifically the message broker, Apache Kafka in this case.



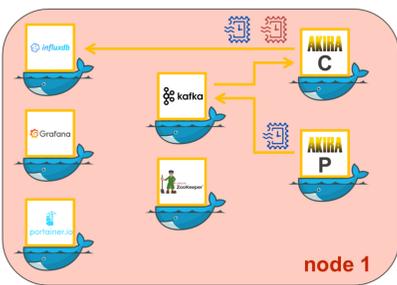
O1. Measure and compare the performance of the message broker and network while scaling up the volume of messages processed

CHALLENGES

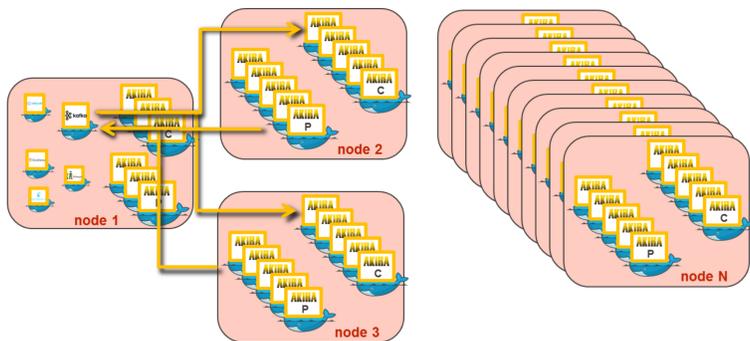
Several challenges were faced:

- Learn how to use Grid5000 and its tool
- Get familiarized with Docker Swarm
- Improve NTP accuracy by using PTPd
- Obtain a suitable approach to reproduce the experiment

DEMO SETUP

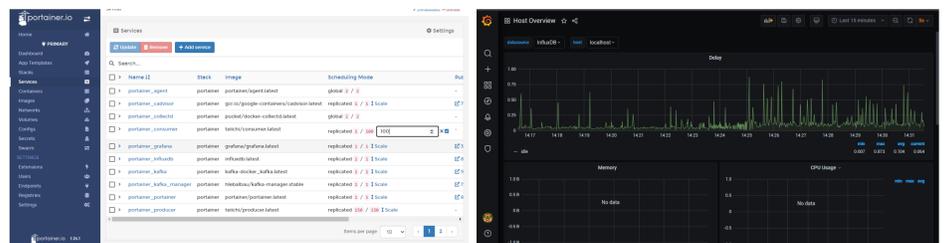


1. Deploy the stack in the swarm
2. Add more nodes
3. Scale containers / services



RESULTS

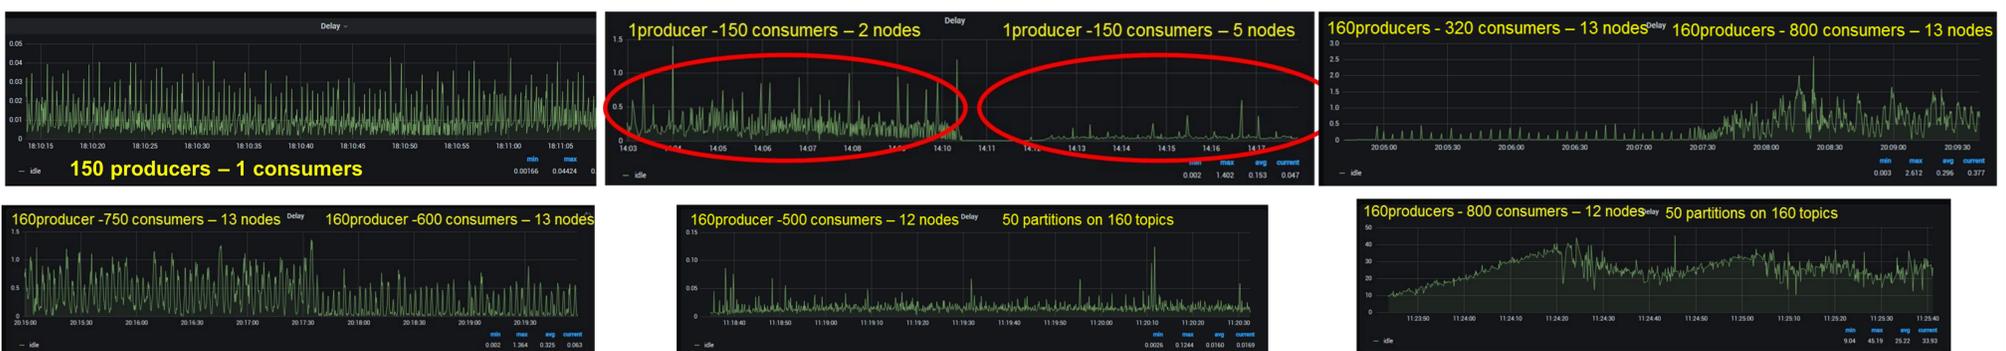
To operate the experiment and get the results, a combination of several tools was used, as Portainer, Influxdb and Grafana



To handle the nodes in Grid5000, some scripting based on OAR was required

```
malvarez@fsophia:~$ ./register and get_ip.sh "/worker_node.sh
docker_token_1598537715209158862 172.16.132.9
Reserving node..
Reserved with IP: 172.16.132.26
```

MORE RESULTS



With the experiment setup shpwn, it has been possible to test different escenarios for different configurations, varying the number of producers, consumers and the number of available nodes. Graphs show the delay time of a message from a producer to a consumer.

CONCLUSIONS

- Docker swarm fits perfectly for managing the scalability of our platform
- Kafka is a very fast and reliable server for our purpose
- Increasing the number of producers does not affect the performance of a server.
- For a single Kafka server, 500-600 producers are the limit for a proper delivery time.
- Fine tuning of the server is required to increase performance and a Kafka servers cluster maybe required.

POST MORTEM

- Having such amount of resources ready to be used is the most appreciated value of Fed4FIRE. Also the good documentation to get hands on quickly is a plus.
- In Stage 2 we will take benefit of different testbeds and will be able to compare different philosophies as in Tengu, GPU Lab and Grid5000.
- Testing Kafka clusters will also be conducted on Grid5000.