CONTAINER-BASED ARCHITECTURE PERFORMANCE ANALYSIS

Pierre-Frédérick DENYS
Monday 6 May 2019

Agenda

- Usage of containers in production
- Monitor a container-based cluster
- Collect traces from a cluster
- Visualization and analysis of traces

CONTAINERS IN PRODUCTION

200 000 containers To run 700 applications

Average lifetime

10% Less than 10 sec

30% Between 5 and 10 minutes

Top techno running in docker :

How it works ?

CONTAINERS ORCHESTRATION

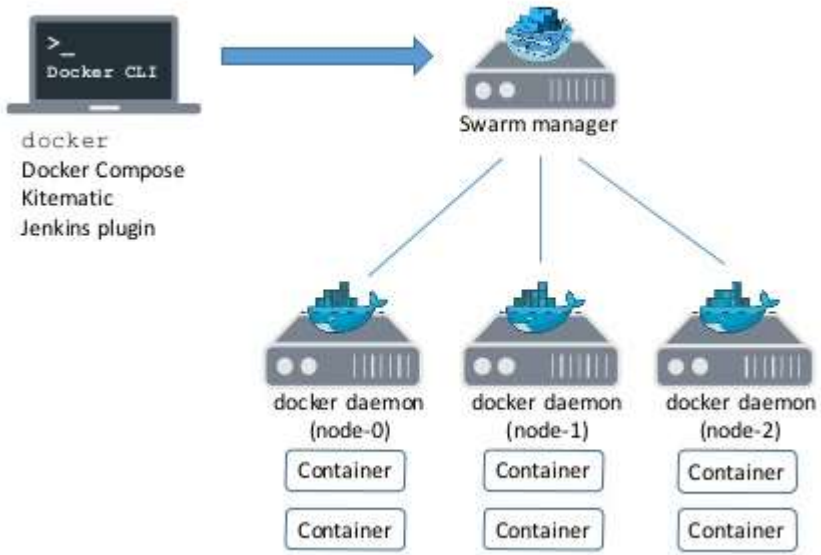

CONTAINERS LIFETIME

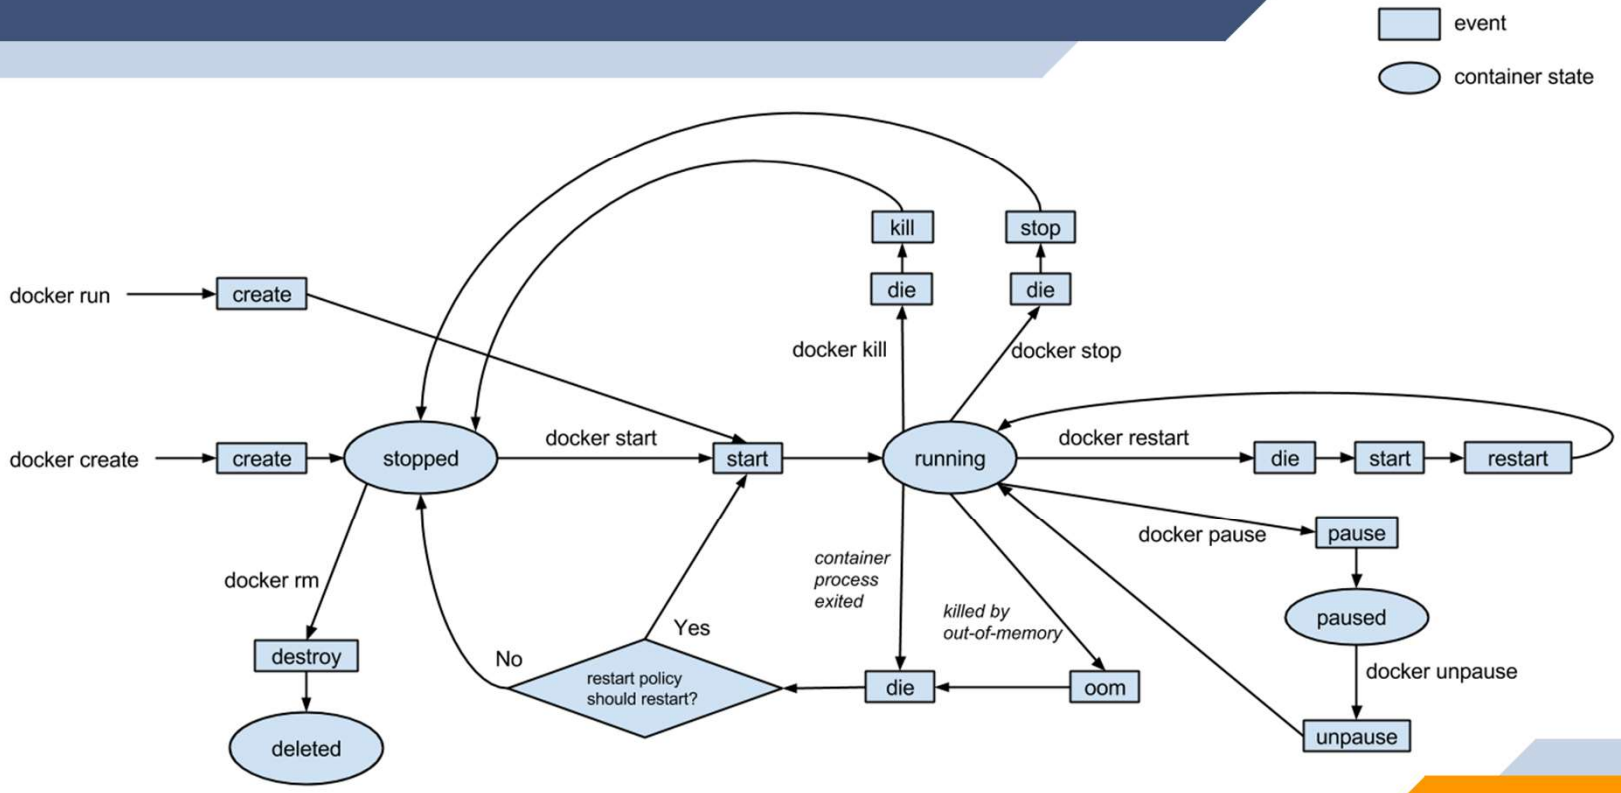

How to monitor these systems ?

PERFORMANCE ANALYSIS SOLUTIONS

- Docker stats
- Cadvisor
- Promotheus
- Heapster (kubernetes solution)
- Elk stack

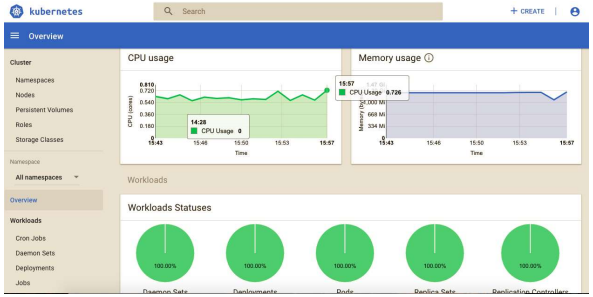

Prometheus

cAdvisor

These tools are high level analysis and based on collected metrics

“ How to solve complex performance problems or random failure?

CONTAINER AND KERNEL

Just a simple process on the host machine => kernel is shared

Cgroups and namespaces are the basis of lightweight virtualization

Cgroups : limit how much you can use
Namespaces : limit what you can see

Thanks to cgroups and namespaces, we have for each containers :

- CPU usage
- Memory usage
- I/O usage per device
- Network usage (cooperation with iptables and tc)

How to collect data ?

DATA COLLECTION

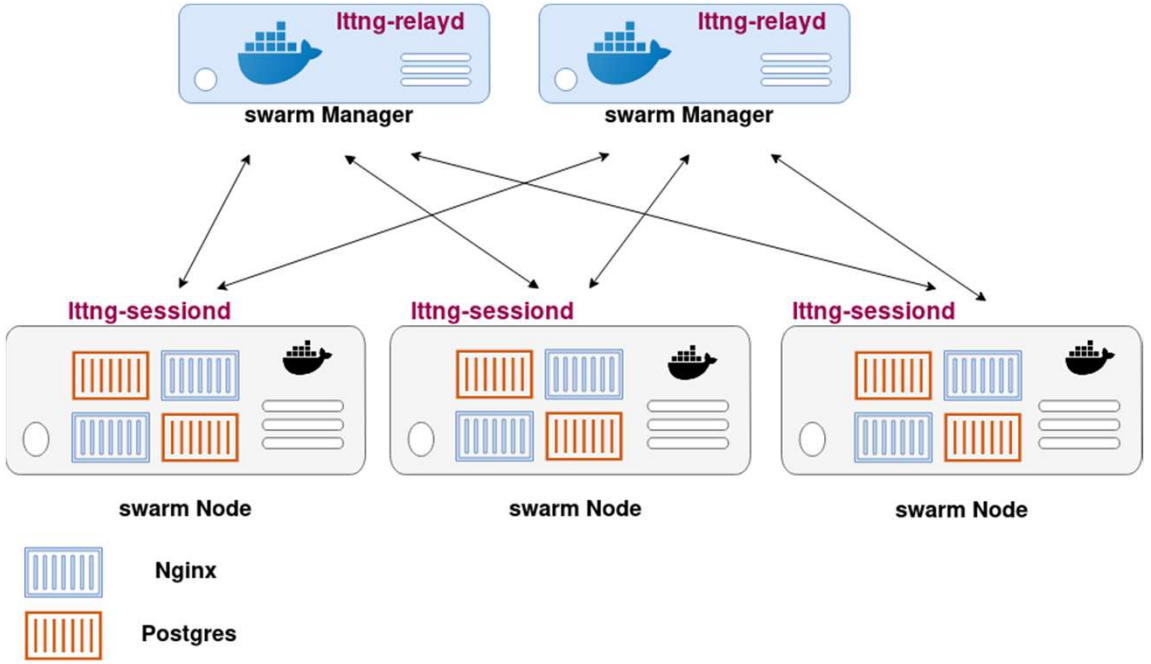

How to add container support to LTTng ?

DATA COLLECTION

What is done actually (queued future version) :

- Kernel tracer (lttng-modules) :
- Namespace contexts to classify and filter events in kernel tracer

What i'll work on :

- Add container runtimes metadata in trace context
- Add state change events of containers
- Adapt LTTng to collect data from several hosts (swarm nodes) in a single
- Make a correlation between the data of theses hosts to analyse a single service

DEVELOPMENT PROCESS

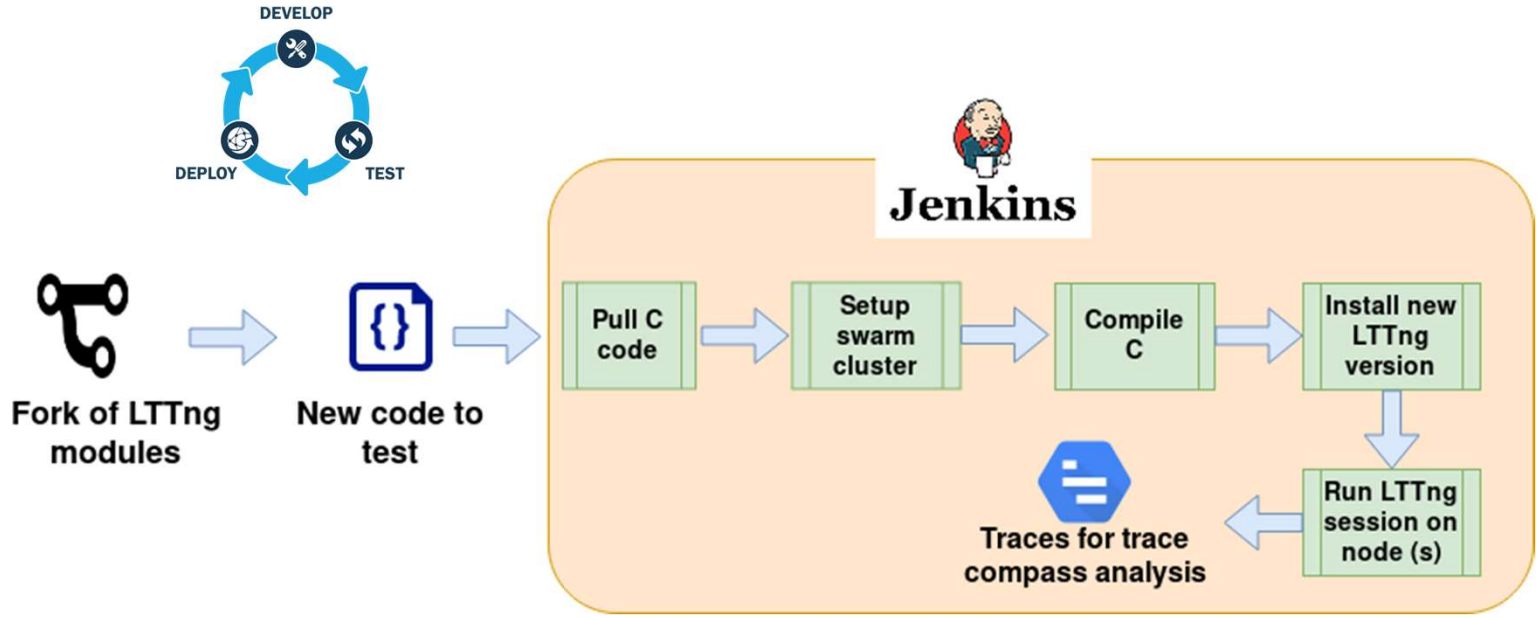

How to analyse and visualize data ?

DATA VISUALIZATION

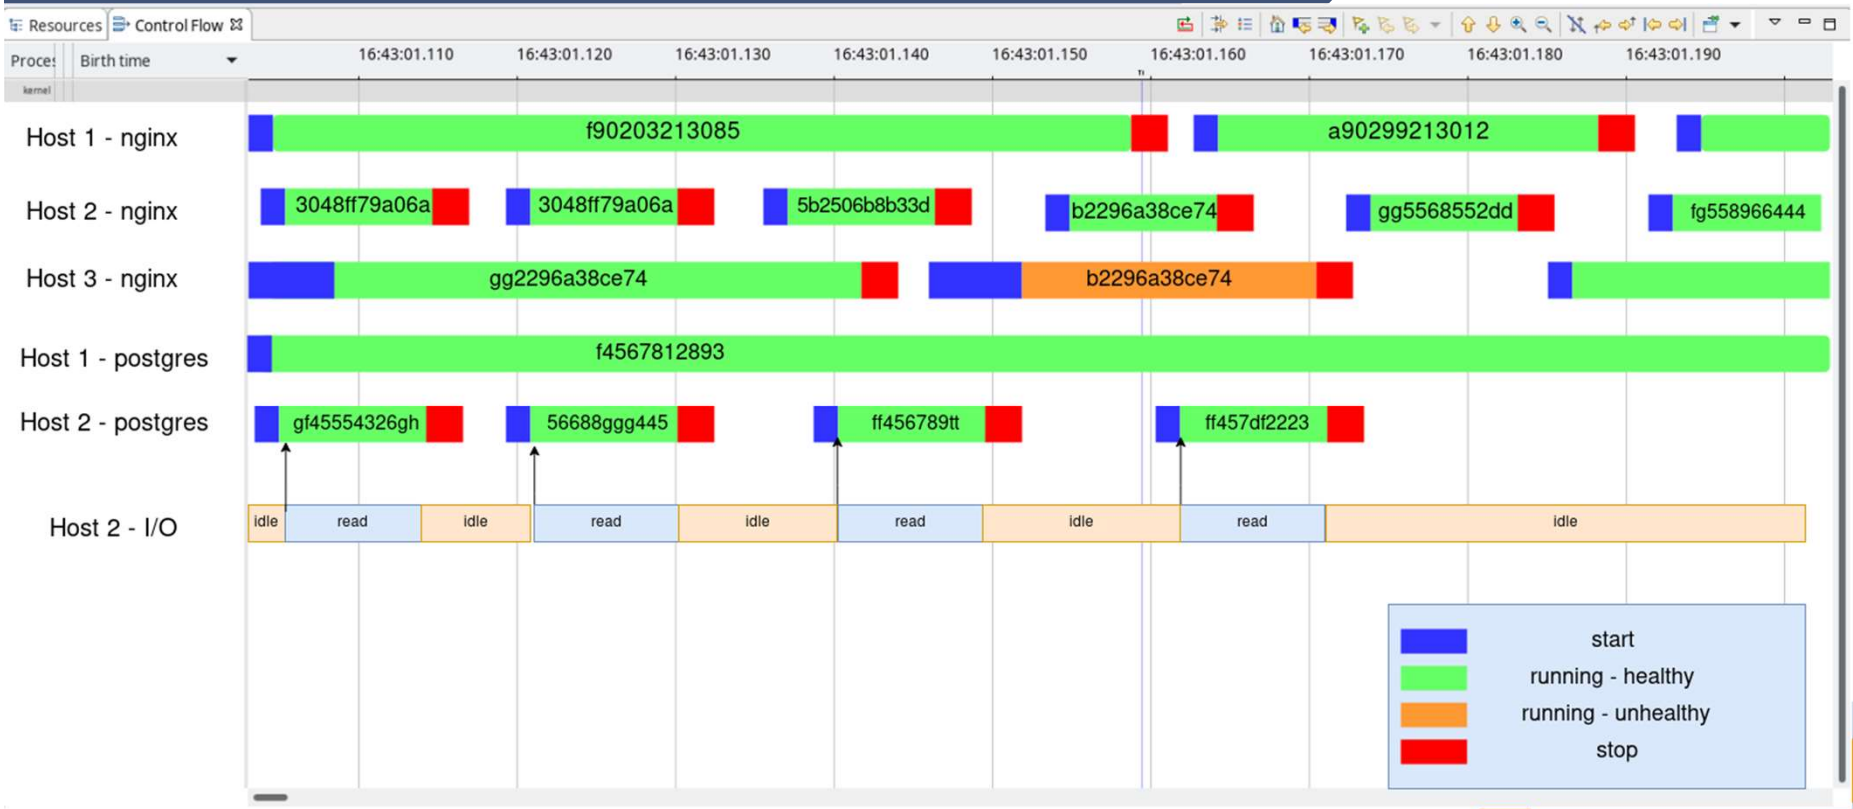

DATA VISUALIZATION

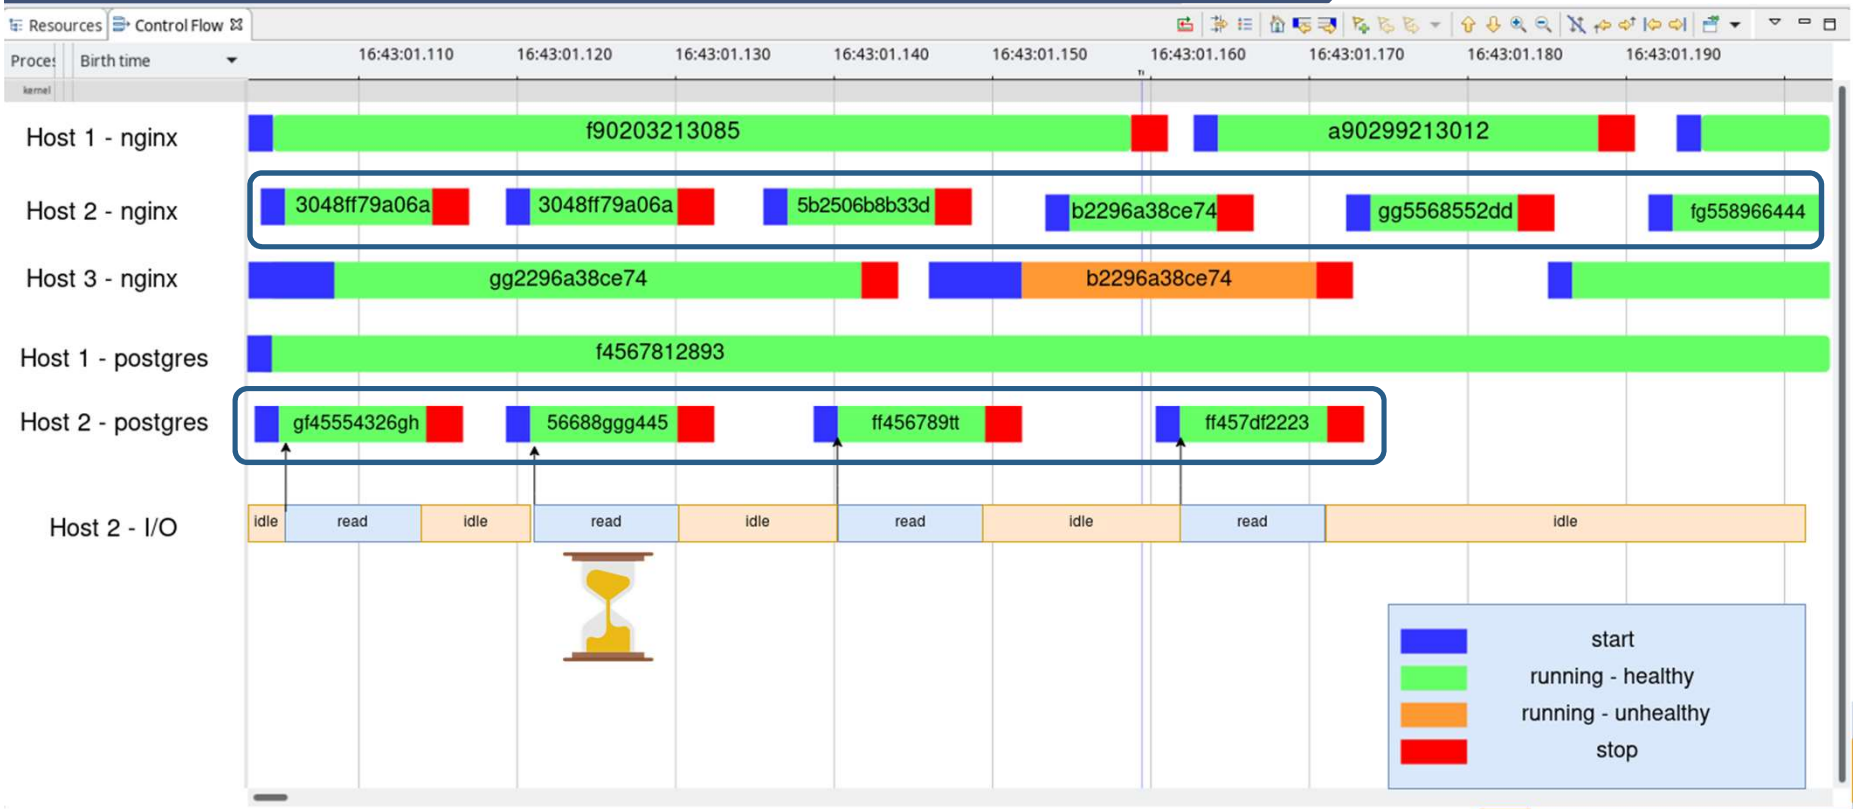

DATA VISUALIZATION

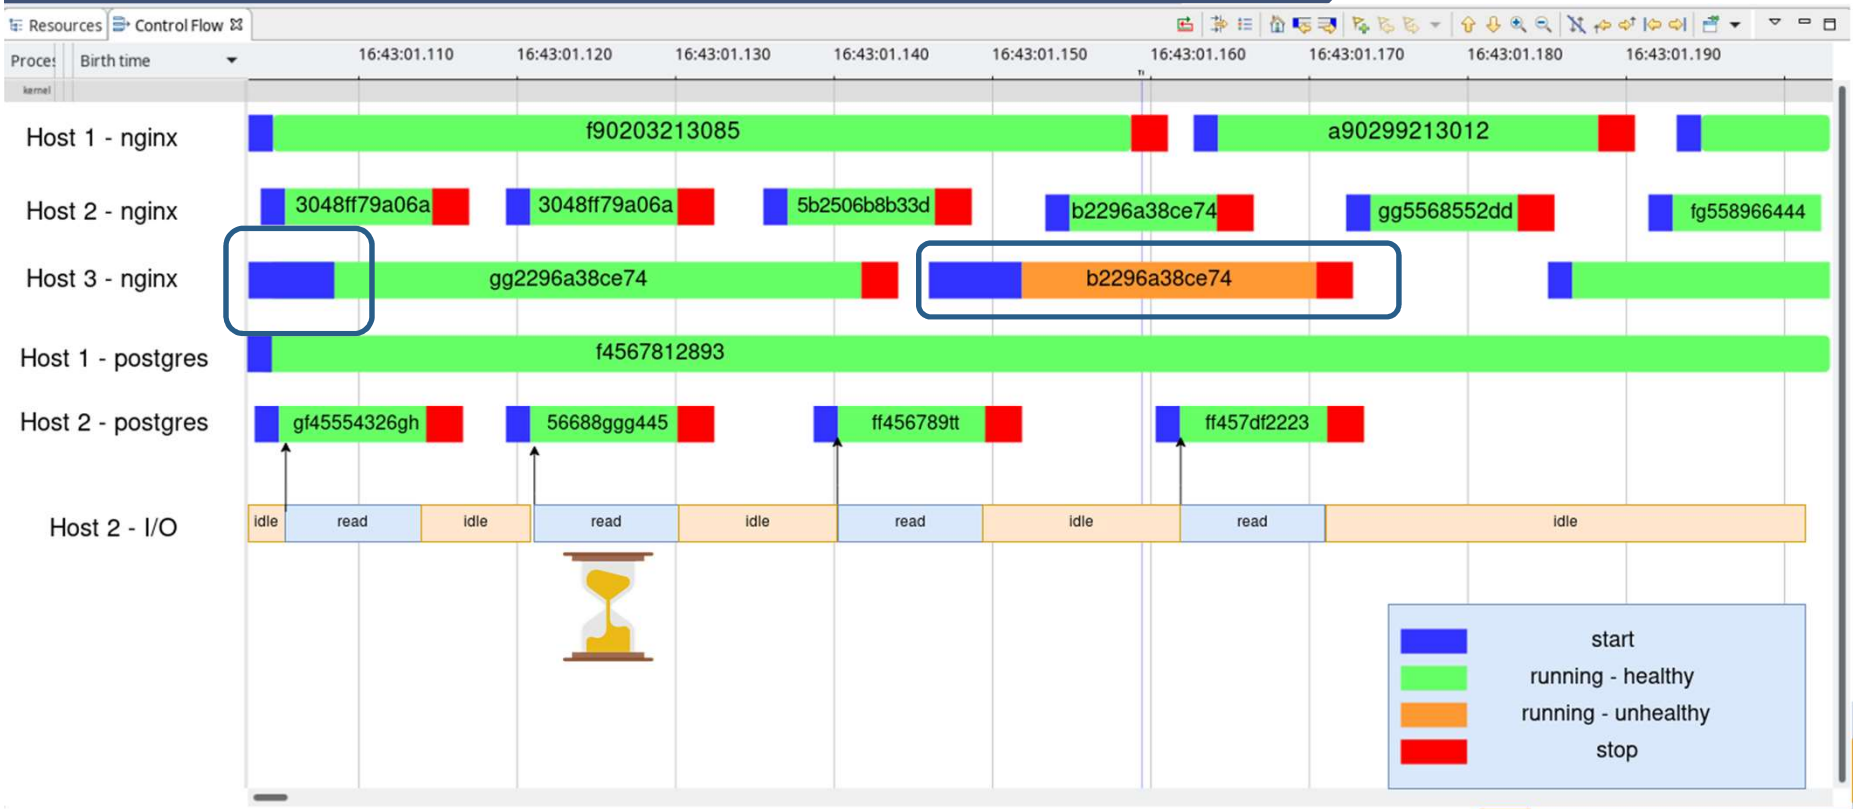


Thank you for
listening !