Monitoring ALICE Analysis Facility infrastructure

challenges and solutions

Pintér Ádám
Wigner RCP – Datacenter

Short introduction

- Started in 2012, CERN Tier-0 for 7y
- 4 + 1 rooms (~436 racks)
- WSCLAB (AF, Tier-2, GPULAB, Virgo/LIGO)
- ALICE Analysis Facility
- Other projects in our datacenter (ELKH Cloud, ELKH ARP)

Capacity

- ALICE Analysis Facility
- Start: November of 2020
- * 3700 vCPU
- * 7.5 TB RAM
- * 2.3 PB storage (raw)
- OS layer: datacenter
- Above: Gergely Barnaföldi, Gábor Bíró & their team

Hardware aging

- Monitoring is an important tool to keep infrastructure healthy and to prolong cluster lifetime.

Defining checks

- Specific environments
- We selected: Sensu
- Server types (worker node, storage)
- Basic checks, advanced and complex
- In almost real time

- keepalive
- check-cpu, check-ram, swap-memory, temperature
- RAID-root-partition, RAID1-status, smart-short-test, disk-read-write-speed
- interface-status, check-internet-access, firewall-access
- switch-ping, switch-PSU, switch-traffic, master-switch-traffic

Network separation

- Dedicated VLAN network for monitoring
- In order to minimize interference with real worker node traffic

Power usage of cluster

- Power consumption and electricity bills are important factors nowadays
- In order to see the details, we are using a visualization solution for power usage statistics based on each rack

Scaling up

- Scale up the monitoring system efficiently
- We are using automation tools for node preparation and installation (Ansible & MAAS)

Event and data history

- Historical data is also very valuable
- We integrated a database solution into our monitoring workflow: InfluxDB

Roadmap for future dev. (I.)

- Continuous disk tests (S.M.A.R.T.)
- Scheduled backup for monitoring data
- Proper alerting based on pre-defined warning and critical levels
- Iterative time-based optimization for running checks
- HTCondor service monitoring

Roadmap for future dev. (II.)

- Monitoring GPU RAM
- GPU-utilization
- GPU card temperature
- Next infrastructure: GPULAB
- Smart alerting for complex cases

Thank you

If you have any questions:

- in person (KFKI campus – building 18)
- pinter.adam@wigner.hu