



Dynamic scheduling using CPU oversubscription on the ALICE Grid

Marta Bertran Ferrer
PhD Student - CERN

On behalf of the ALICE Collaboration

Grid resource usage overview

Optimize the use of resources by taking advantage of the different nature of Grid payloads

- Analysis jobs are I/O intensive and they do not fully use the allocated CPU slot
- **Idle portions of CPU** can be gathered and used to execute computing intensive jobs (MonteCarlo)
 - No added pressure in I/O

Our analysis shows that ~75% of the worker nodes have **spare memory resources** not used by the running jobs

- Base allocation of 2GB RAM/core (up to 8GB/core with SWAP)

Whole-node allocation and partitioning

Multicore jobs can run in two configurations

- **Whole node:** JAliEn job wrapper manages resources and splits them into job execution slots
- **Slots of N cores:** Jobs are assigned a set of resources by batch
 - This might be further enforced by some resource constraining mechanism
 - These slots can be further partitioned to run finer-grained jobs as in the whole node case

We are particularly interested in the a **whole node** configuration - allows for increased flexibility and optimal resource management

- Boosting fine-tuned late binding job allocation

Whole-node allocation and partitioning

In whole node scheduling scenarios we want to use the available resources in the **most efficient way**

- Memory is the biggest constraint during job execution
- Compute **memory per CPU core ratio** and see if they can be oversubscribed
- Node's *Idle* CPU above a threshold → It has room for an extra job
- Assign unused memory to new jobs and oversubscribe CPU cores

Making use of the internal machine and job resource monitoring

Analysis of current Grid worker nodes usage levels

Analysis of the **current usage levels of the Grid worker nodes** during **72 hours**

Goal - to have an accurate picture of the **feasibility of adopting oversubscription policies** for the hosts used for **whole node** submission

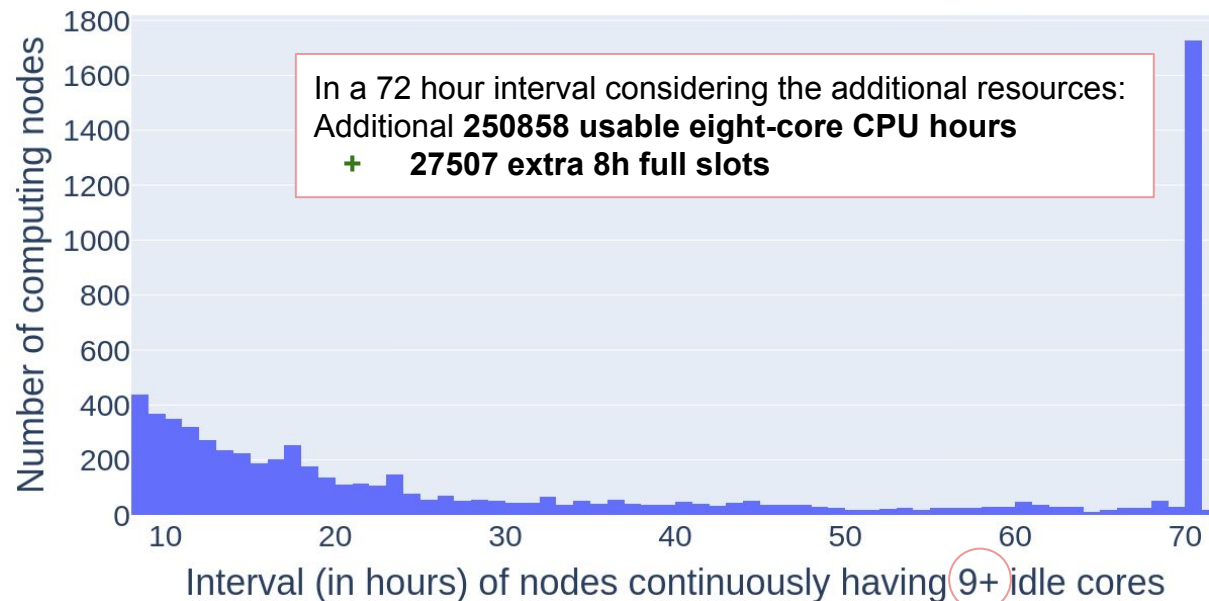
All hosts reporting values to ALICE monitoring system MonaLisa have been included

- To take into account that they might be running other workloads in parallel for which we do not have any knowledge → ALICE workflows are also heterogeneous

Continuous sampling of the **amount of idle cores** to evaluate the potential additional slots

- Set **grace interval** for deciding **when to start** a new payload (waiting for the amount idle cores to stabilize) and for deciding **when to kill** a payload (when machine becomes saturated)

Eight-core usable resources in oversubscribed mode

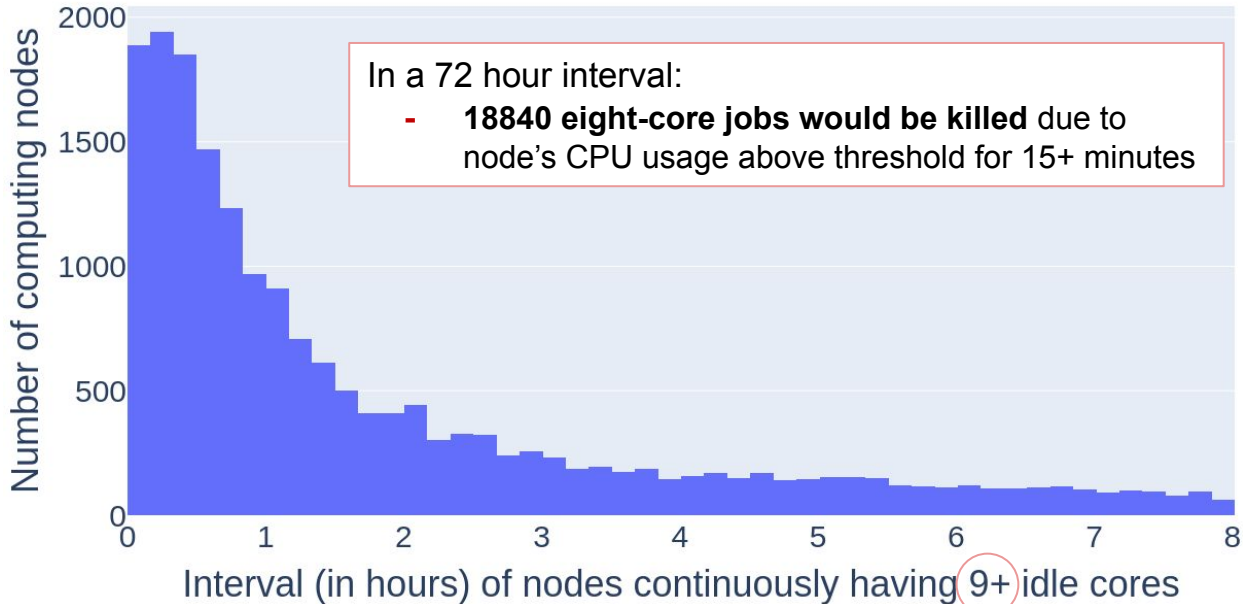


Usable eight-core CPU hours:
Idle time intervals lasting 8+ hours

Total eight-core CPU hours of the surveyed hosts in 72h: 2466117 hours => Additional usable CPU hours (from plot on left) are a **10.2% of the total**

8 core job execution + 1 core left idle for system operations

Caveat: Potential eight-core jobs preemption

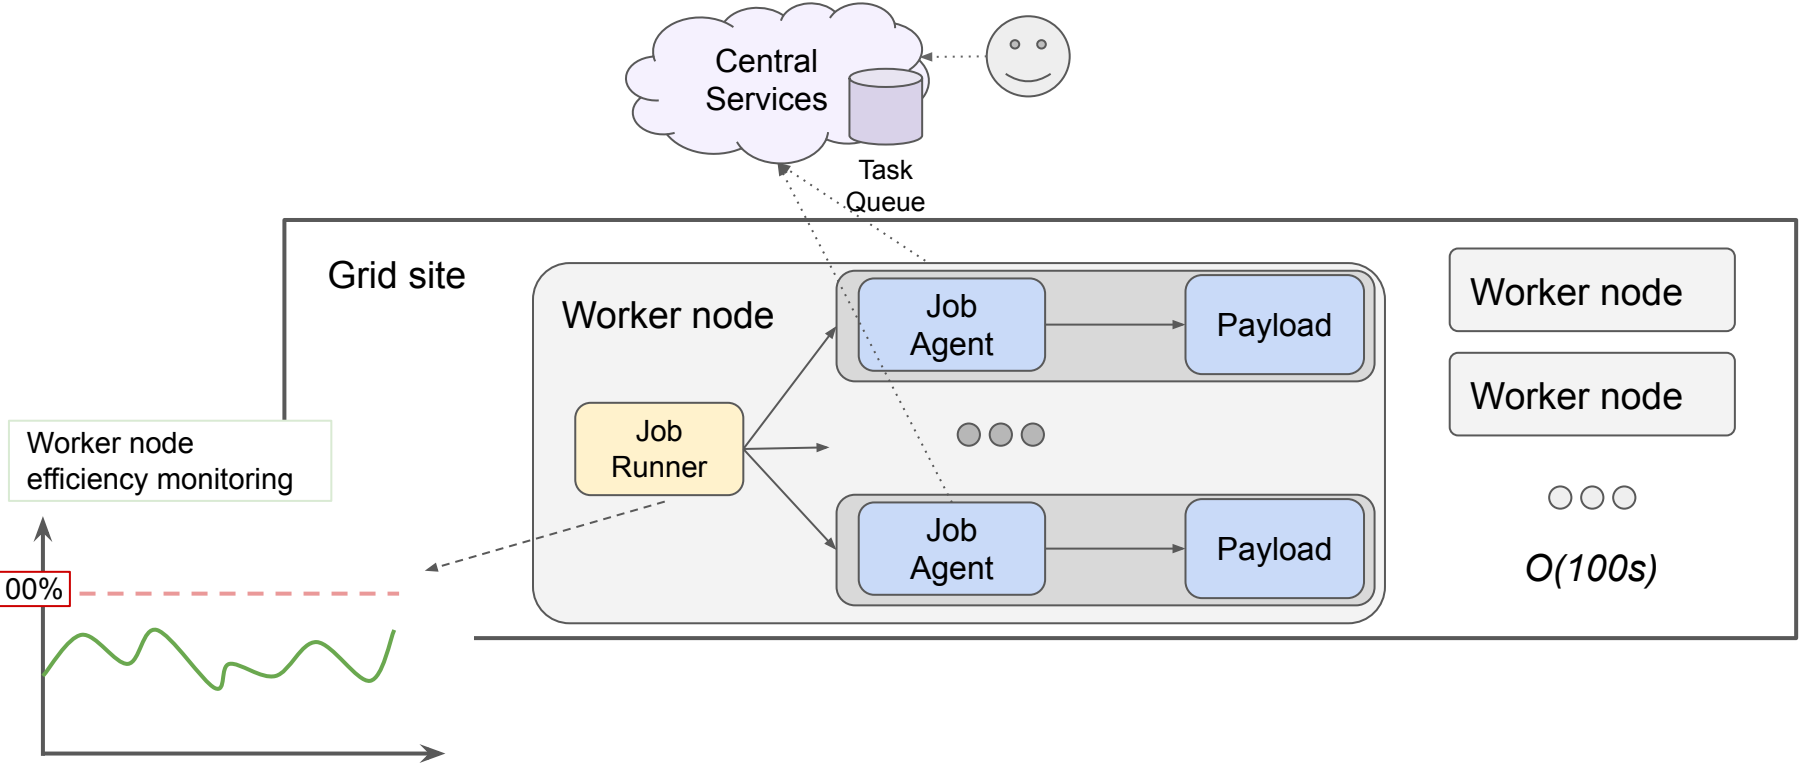


Monte-Carlo workflows that run in **opportunistic fashion**

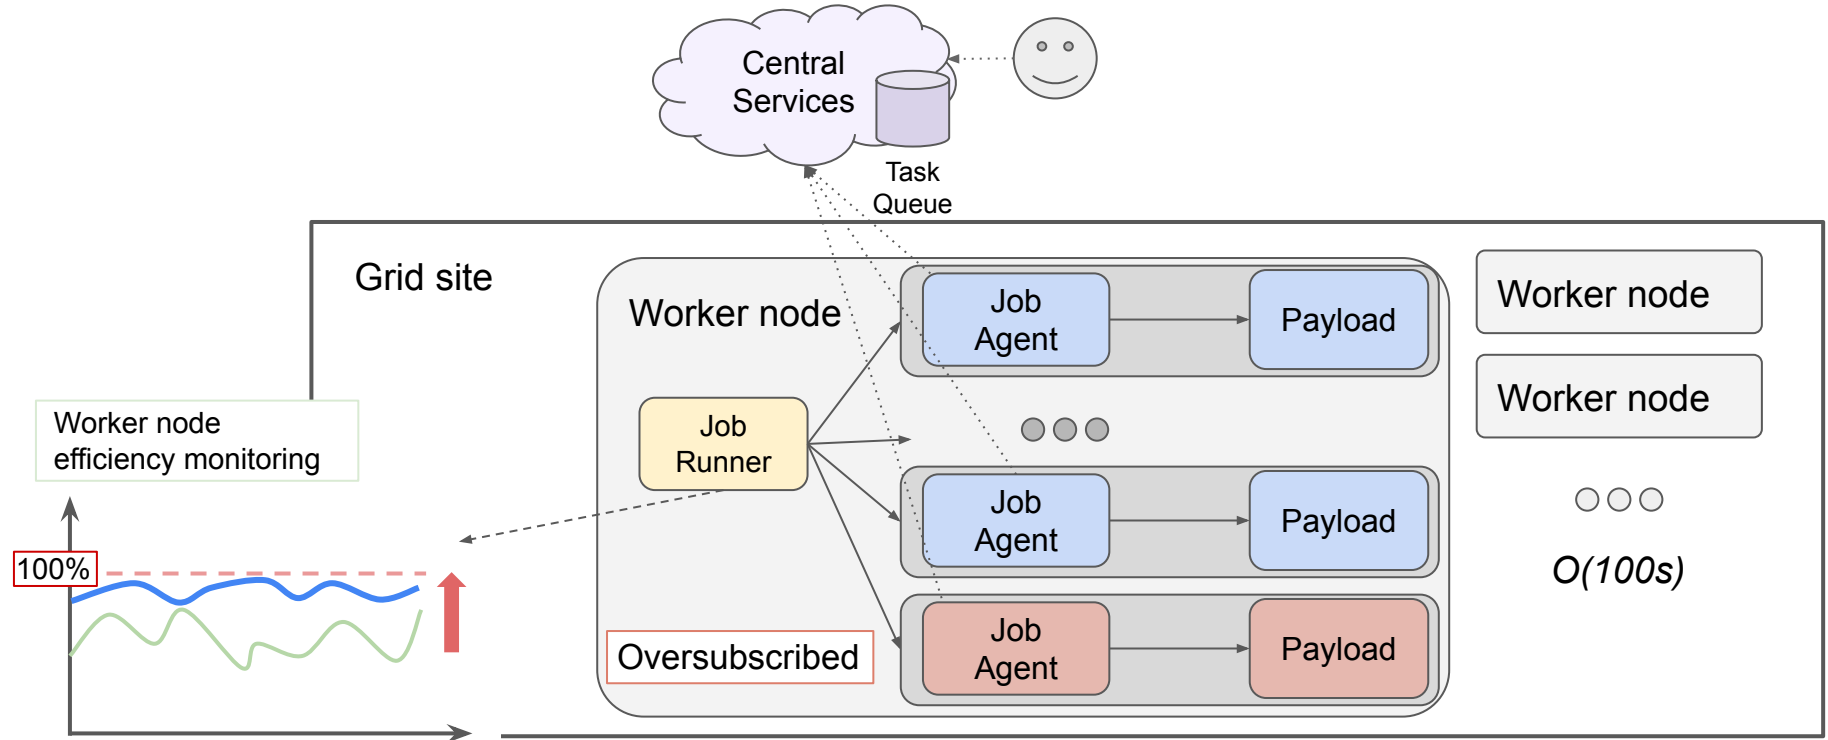
- Find a balance between preemption rate and pressure in database

8 core job execution + 1 core left idle for system operations

JAliEn in oversubscription regime



JAliEn in oversubscription regime



Oversubscription implementation

The amount of CPUs available for oversubscription will depend on the node's memory

- Minimum between $[\text{RAM}/2, (\text{RAM}+\text{SWAP})/8]$

Jobs to execute under oversubscription regime are CPU intensive (minimal I/O)

- Target jobs are MonteCarlo with a number of cores that fits the allocated amount of CPUs in the oversubscription pool

Oversubscribed jobs are executed re-niced → **Lower priority**

Preemption of oversubscribed jobs if machine gets saturated

- Jobs are rescheduled without penalty
- Priority to continue executing long-lived jobs

Deployment in production environment


- The oversubscription-enabled framework has been deployed on two test machines
 - Oversubscribed cores limited by SWAP in both cases:
 - Machine 1: 32 cores Intel Xeon Gold 6244 at 3.60GHz. 11 oversubscribed cores
 - Machine 2: 64 cores Intel Xeon Gold 6226R at 2.90GHz. 3 oversubscribed cores
- Deployment of 24h slots with fully saturated machines executing equal single-core analysis jobs
 - Oversubscribed jobs are single-core CPU-intensive MonteCarlo simulations

Studied metrics


- Job efficiencies
 - Regular job efficiency $\rightarrow \frac{\Sigma \text{CPU time regular jobs}}{\Sigma \text{walltime regular jobs}}$
 - Extra job efficiency (MonteCarlo) $\rightarrow \frac{\Sigma \text{CPU time extra jobs}}{\Sigma \text{walltime extra jobs}}$
- Machine utilisation efficiency
 - $\frac{\Sigma \text{CPU time all jobs (regular + extra)}}{\text{Slot duration (24h)}}$
 - This is the **metric we aim to improve**

Oversubscription behaviour - Machine 1

CPU usage information



CPU usage information



Evaluation

- Taken for realtime = 24h.
- $\frac{\Sigma \text{CPU time all jobs (regular + extra)}}{\text{Slot duration (24h)}}$
 - $\frac{\text{Utilization efficiency oversubscribed machine 1}}{\text{Utilization efficiency non-oversubscribed machine 1}} = \mathbf{1.28}$
 - $\frac{\text{Utilization efficiency oversubscribed machine 2}}{\text{Utilization efficiency non-oversubscribed machine 2}} = 1.03$

Conclusions

- **Whole node scheduling** allows for optimal management of the available resources
- The increase in machine utilisation compensates the decrease of efficiency of the execution of individual jobs
- Observed increased utilisation efficiency of the hosts is proportional to the amount of extra executed jobs
- The decrease in job efficiency is also compatible within sites considering the amount of oversubscribed tasks
- Running additional jobs with **complementary resource usage patterns** on a worker node has a great potential to compensate for inefficiencies due to the different nature of the payloads (I/O intensive vs CPU intensive)