

– Evolve – Results – in Brief

Deliverable D5.4

First report workflow resource allocation mechanisms and policies.

By: Christos Kozanitis . FORTH

THE PROBLEM

OBSERVATION

• Users tend to overestimate the resources that they reserve

• 3-5x more CPU, RAM than their workloads need

WHY?

- Because they plan for the worst case execution
- Usually, only a fraction of their running time do workloads meet the worst case

RESULT

- Expensive deployments
- Need to buy more hardware for more users/workloads

OUR APPROACH

- Designed for analytics workloads that share resources with HPC workloads
- Dynamic estimation of resources
- Implemented as a Kubernetes scheduler that expects users to enter target performance
- The module monitors performance targets. It increases resources when the observed performance is smaller than target performance, and it decreases resources when the opposite happens. If a node runs out of resources, the module migrates the pod.

Dynamic allocation on Big Data workloads improves collocated HPC runs up to 10%

