

SCHEDULING HYBRID WORKLOADS IN SHARED CLOUD INFRASTRUCTURES

DALIBOR KLUSÁČEK¹ AND GABRIELA PODOLNÍKOVÁ²

¹CESNET a.l.e., Czech Republic — ²Faculty of Informatics, Masaryk University, Czech Republic
klusacek@cesnet.cz, xpodoln@fi.muni.cz

1. MOTIVATION

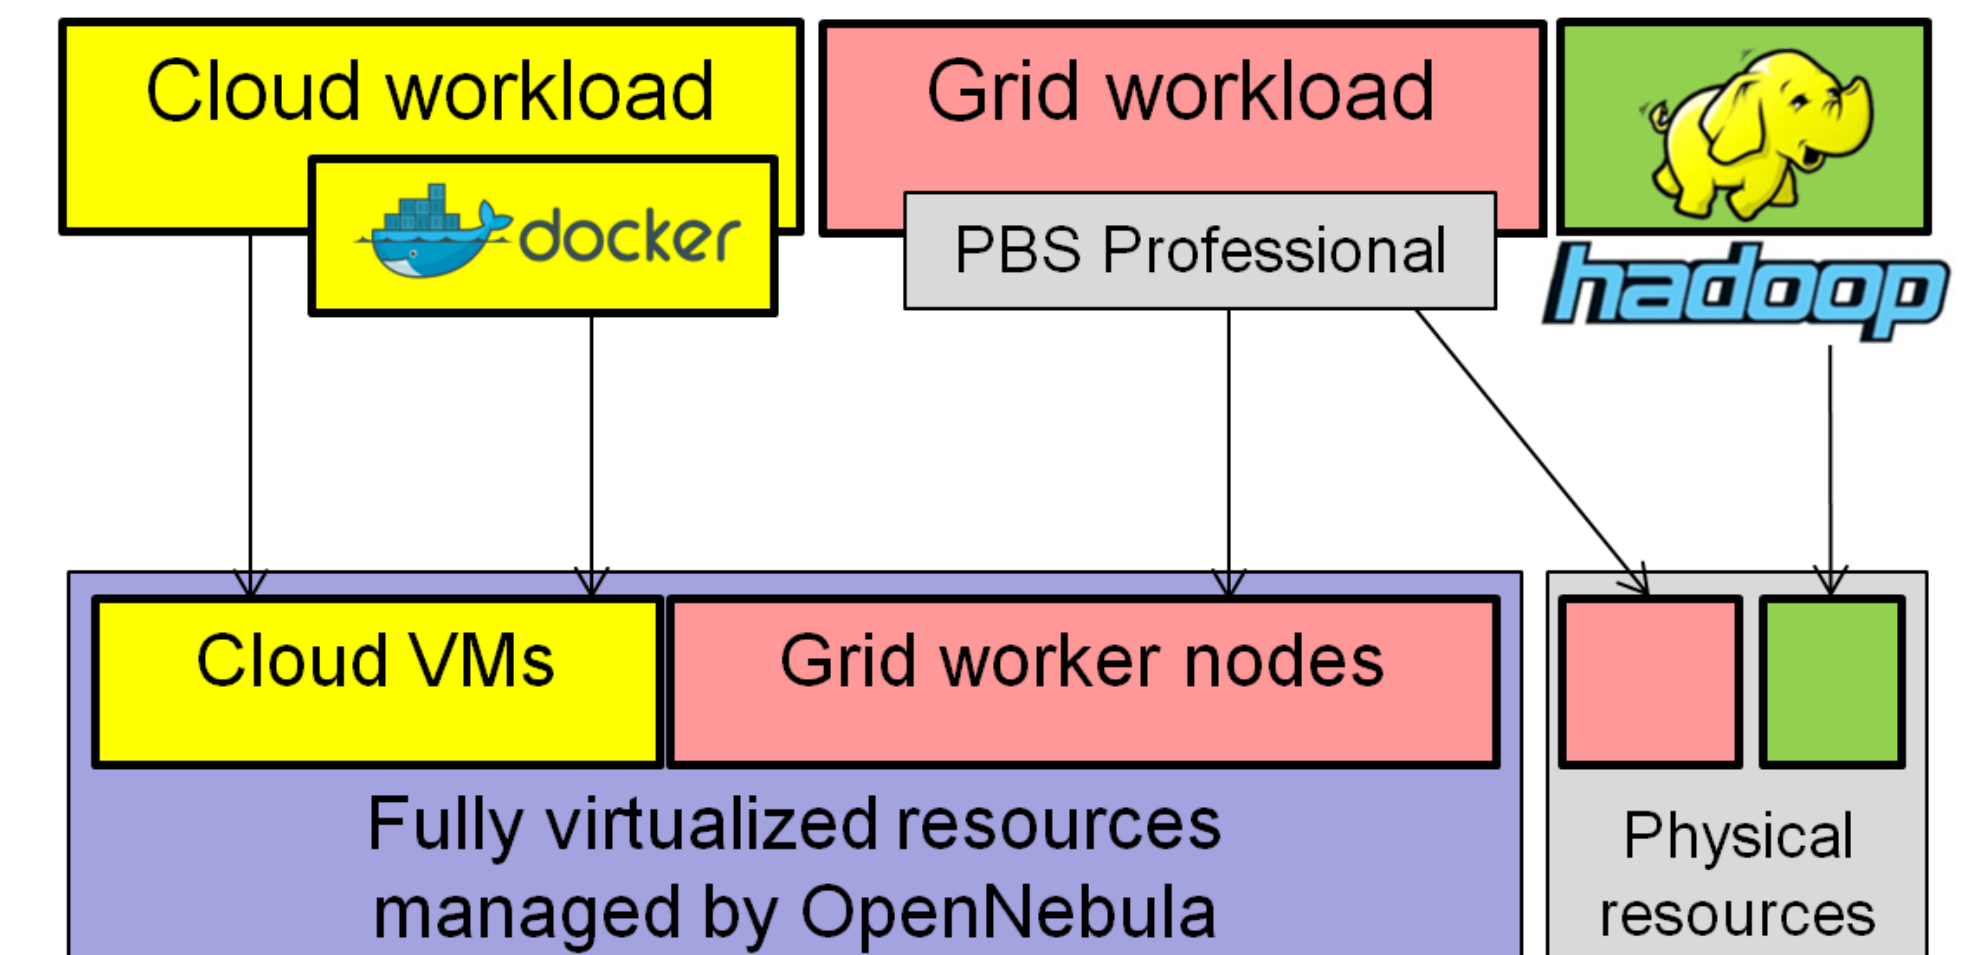
This work describes the recent research concerning **scheduling in shared cloud infrastructures**. Our goal is to automate and improve current status by introducing:

- automated load-balancing
- reclaiming of inactive resources
- advanced fair-sharing mechanisms
- improved VM scheduling policies

2. SHARED CLOUD-BASED INFRASTRUCTURE

MetaCentrum infrastructure:

- infrastructure is mostly virtualized
- currently using OpenNebula platform
- delivering flexible IAAS
- VMs may host grid worker nodes
- PBS-Pro uses grid worker nodes
- load-balancing done “by hand”

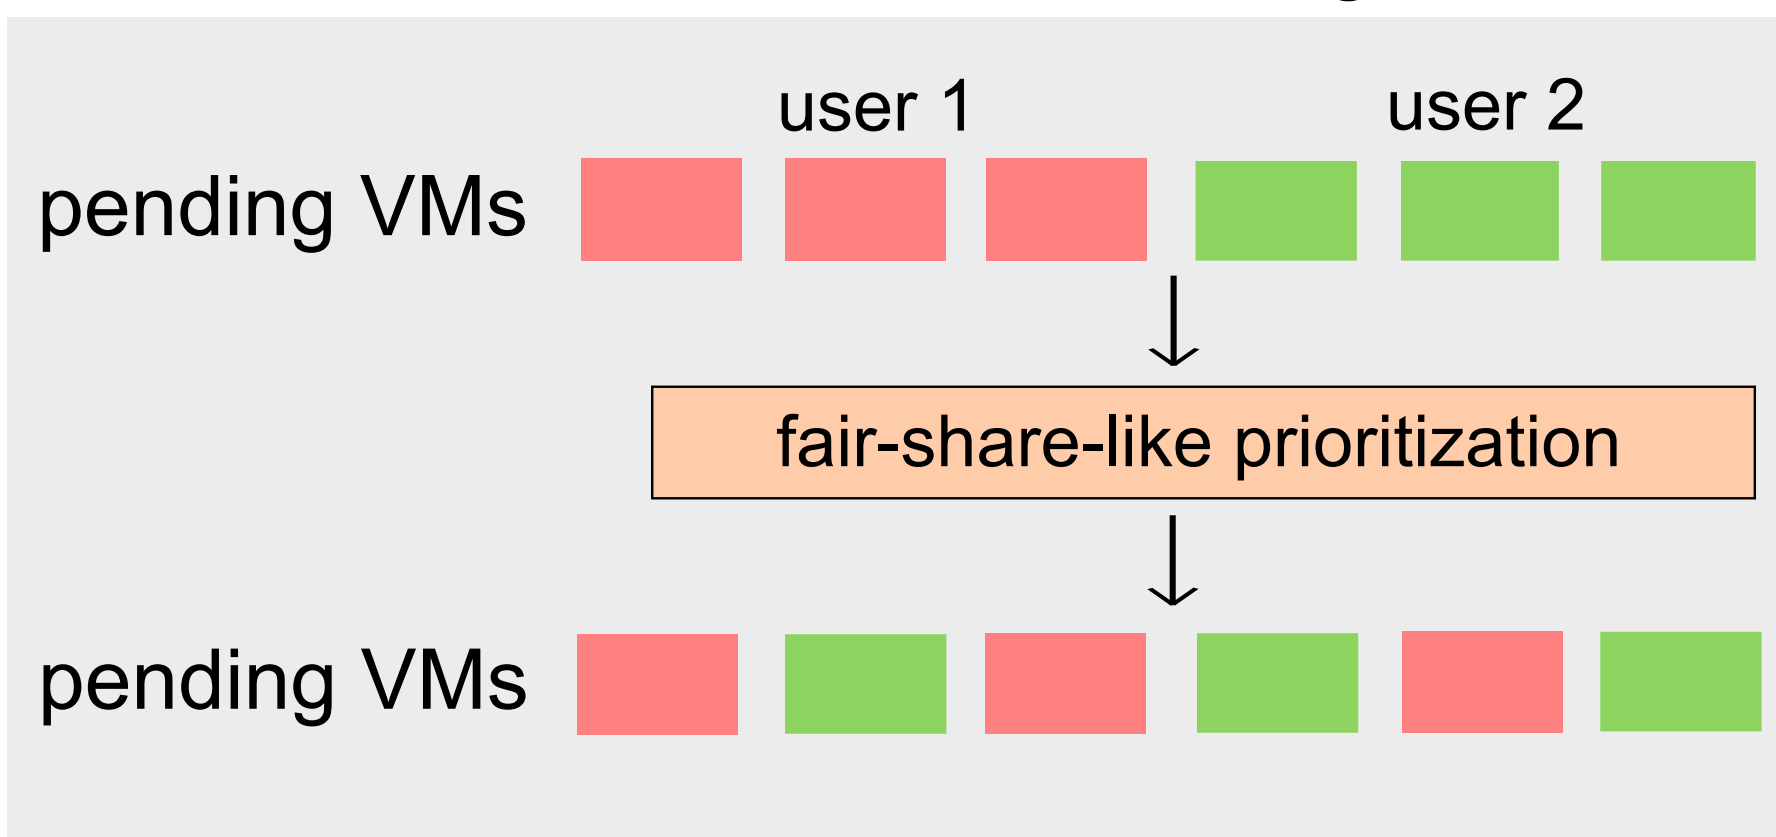


3. NEW VM SCHEDULER

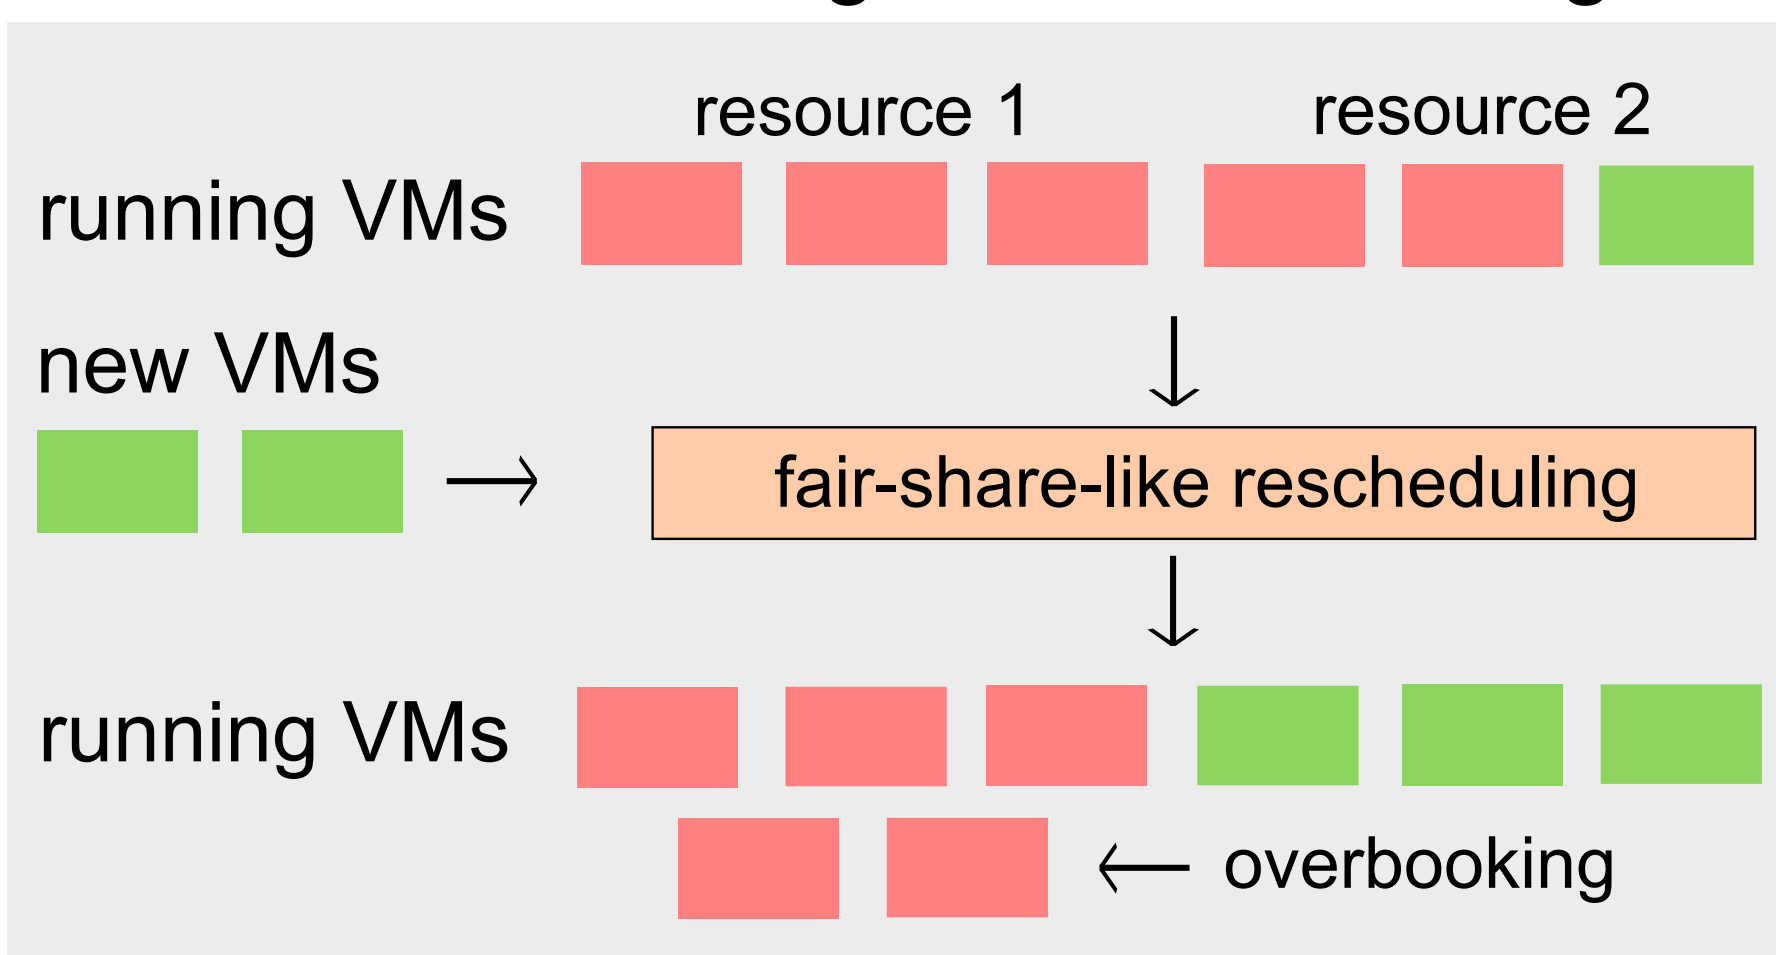
New scheduler for OpenNebula [1, 3]:

- multiple queues for pending VMs
- application of complex policies
- multi-resource aware fair-sharing [2]
- multiple (re)scheduling approaches

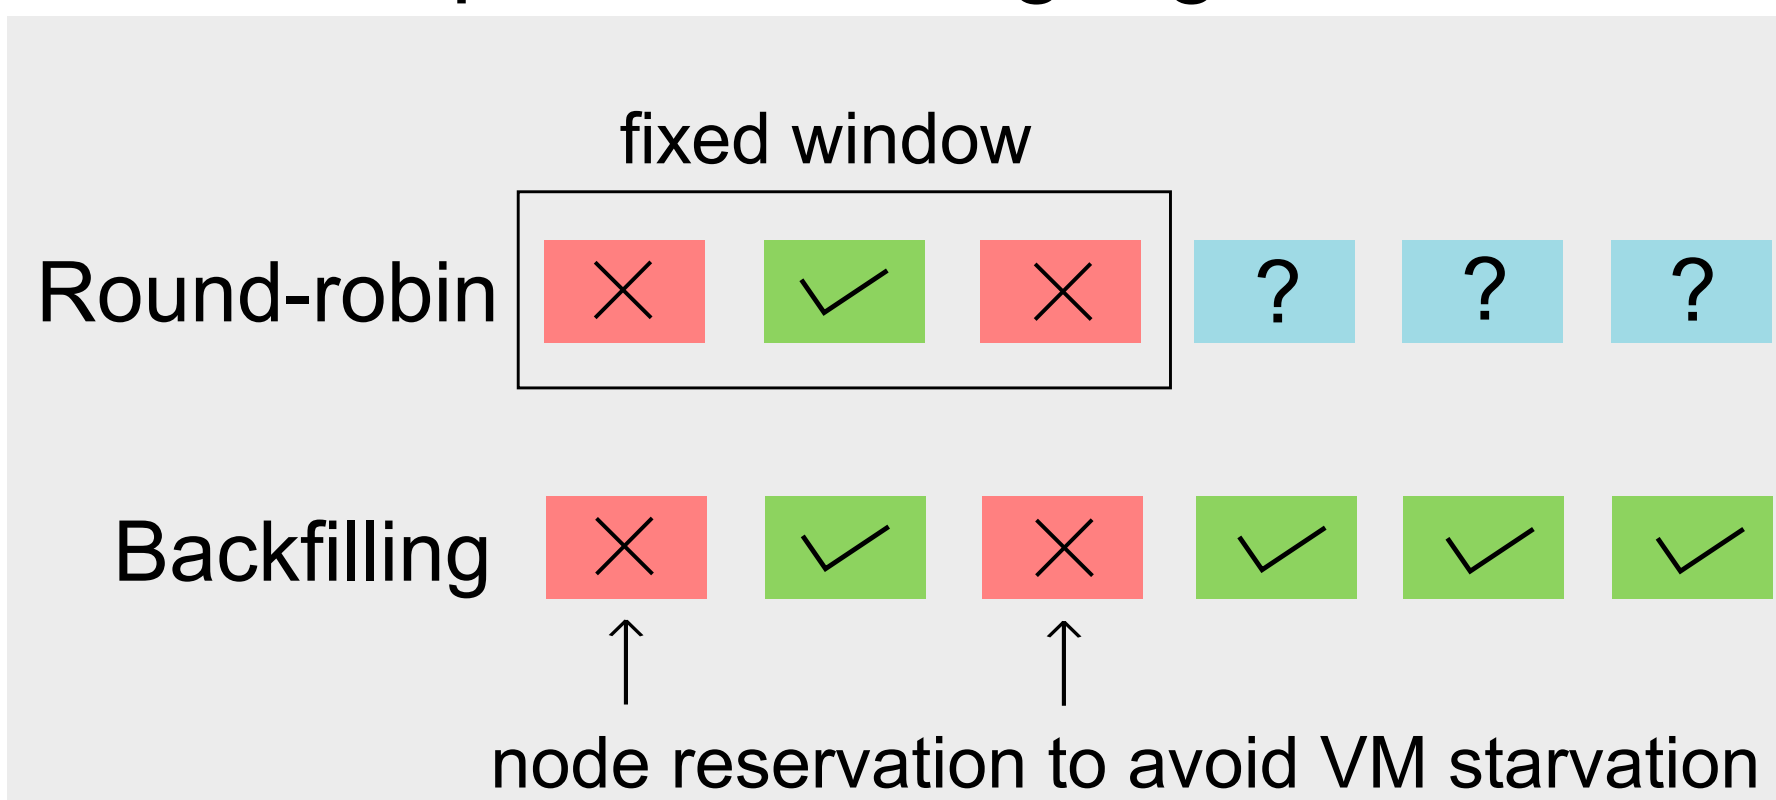
Resource fair-sharing



VM rescheduling and overbooking



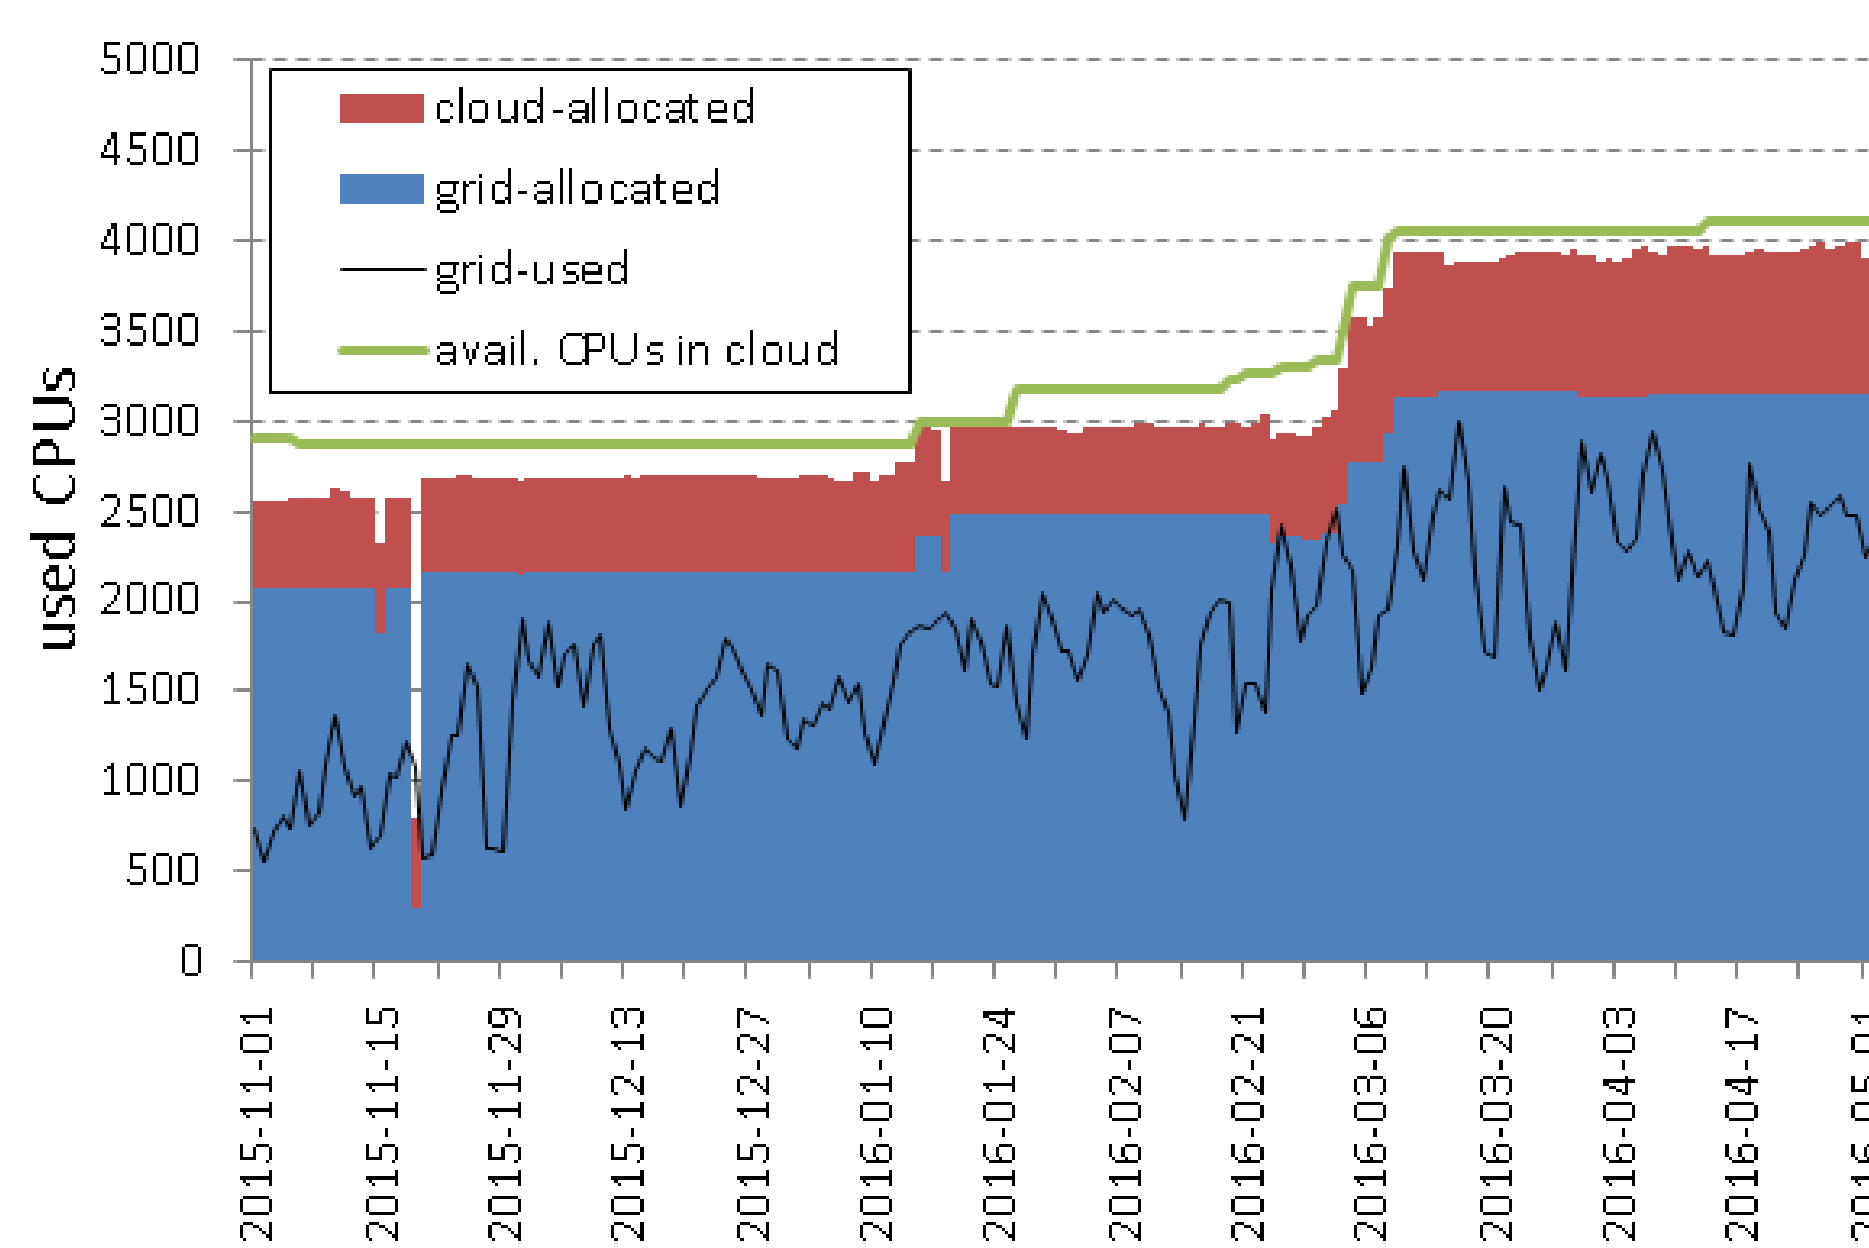
Multiple scheduling algorithms



4. HYBRID WORKLOADS AND GLOBAL SCHEDULING

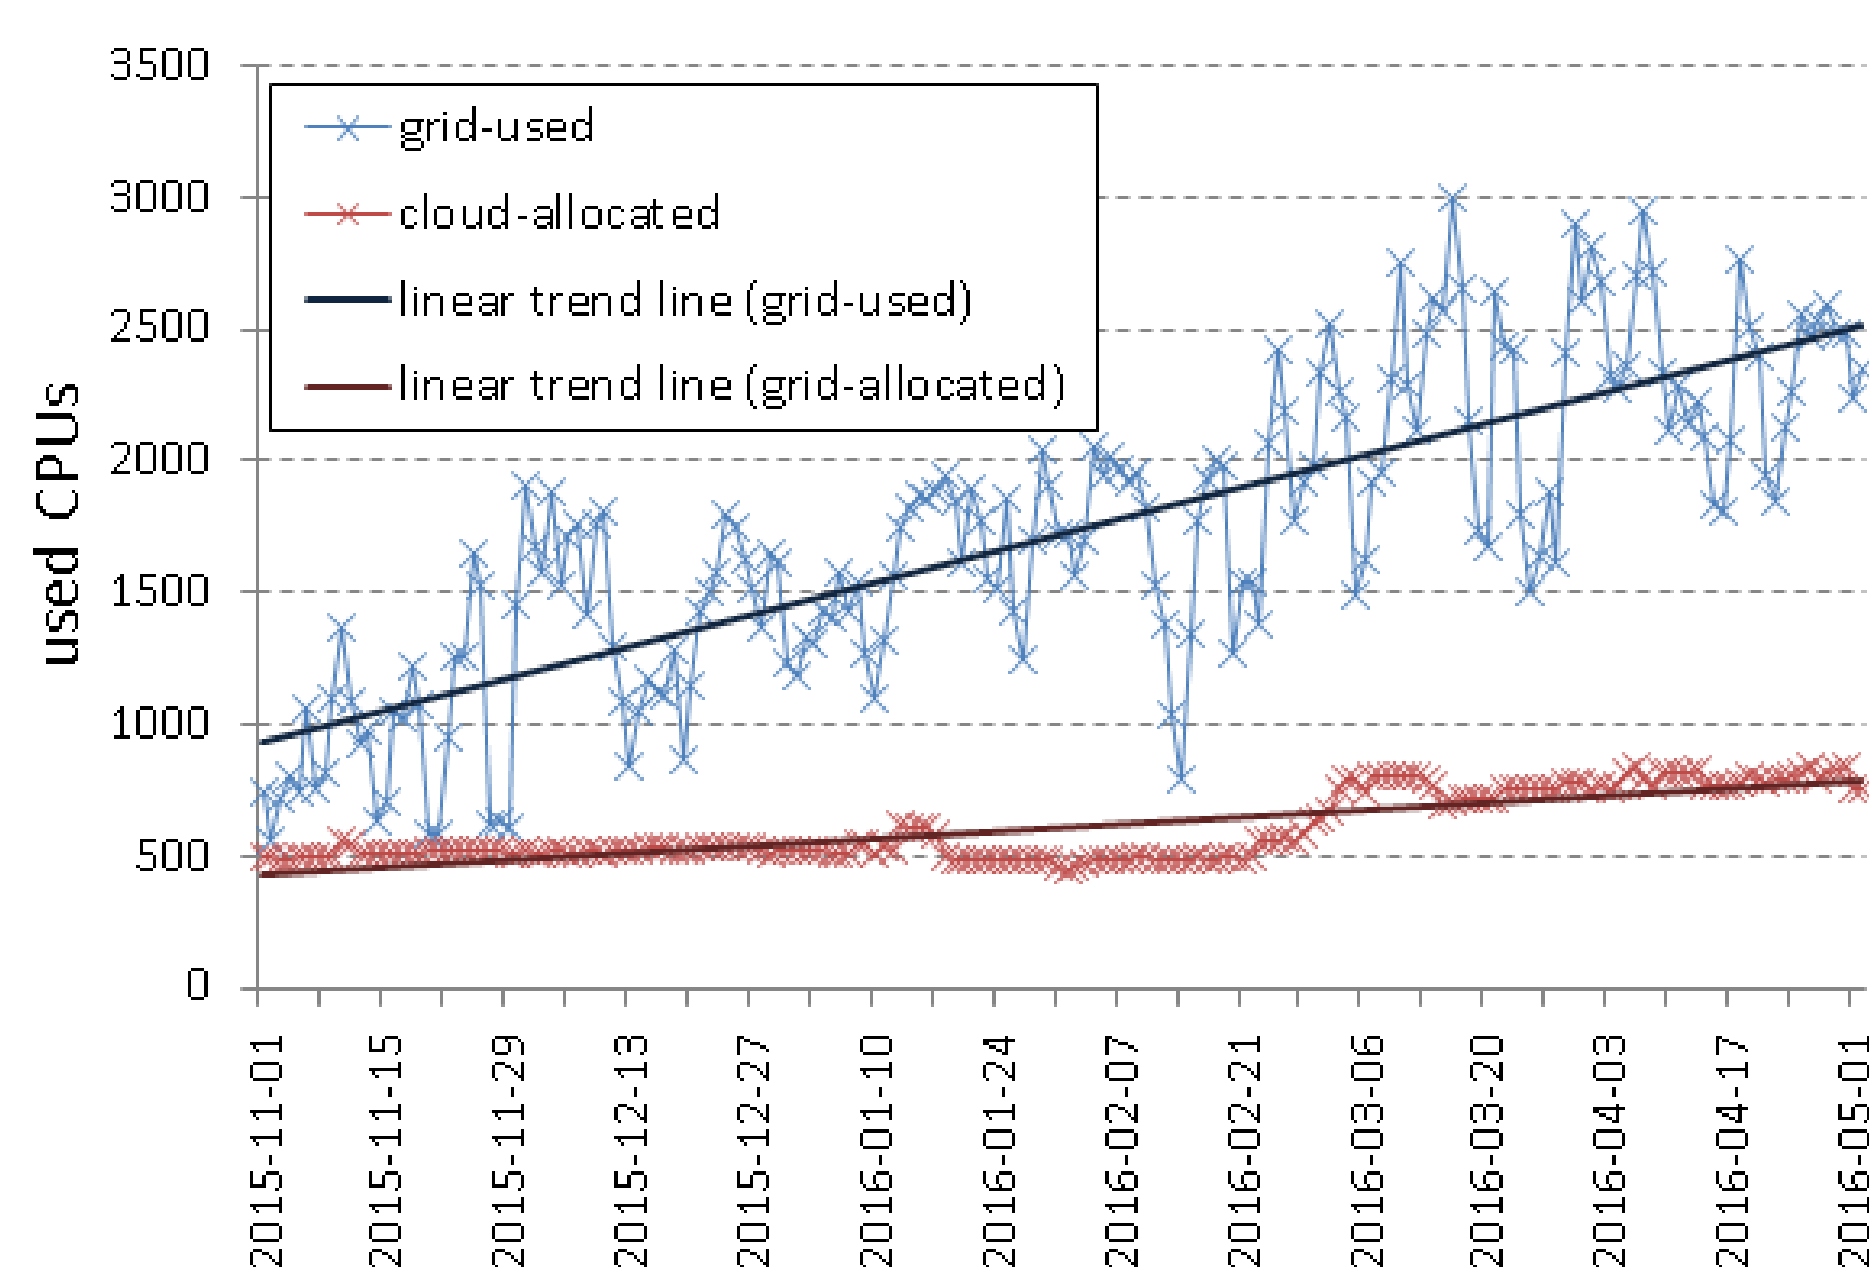
Workload Characteristics:

- hybrid (mixed) workloads
- cloud VMs and grid worker nodes
- grid worker nodes execute “grid jobs”



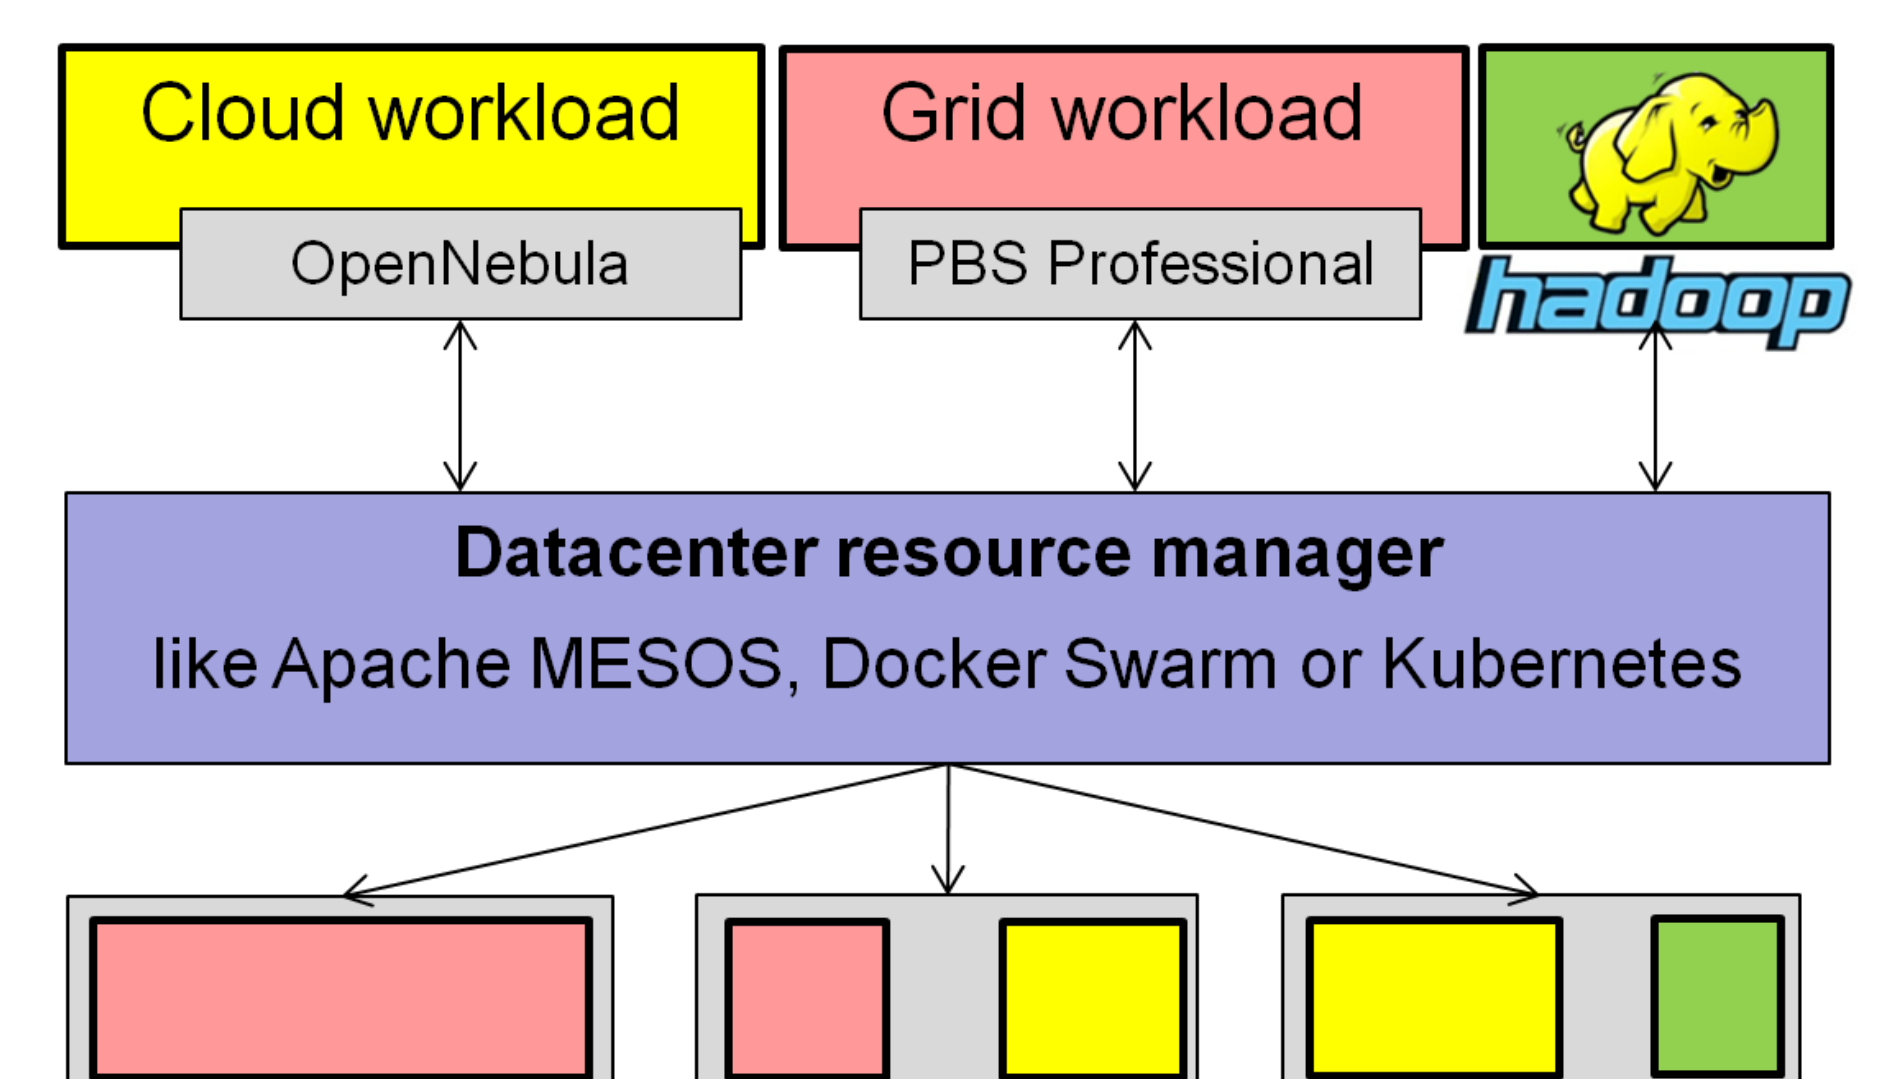
VM life-cycle problems:

- “pure” cloud workload is growing
- pay-per-use model is not used
- very low CPU/RAM reclaiming



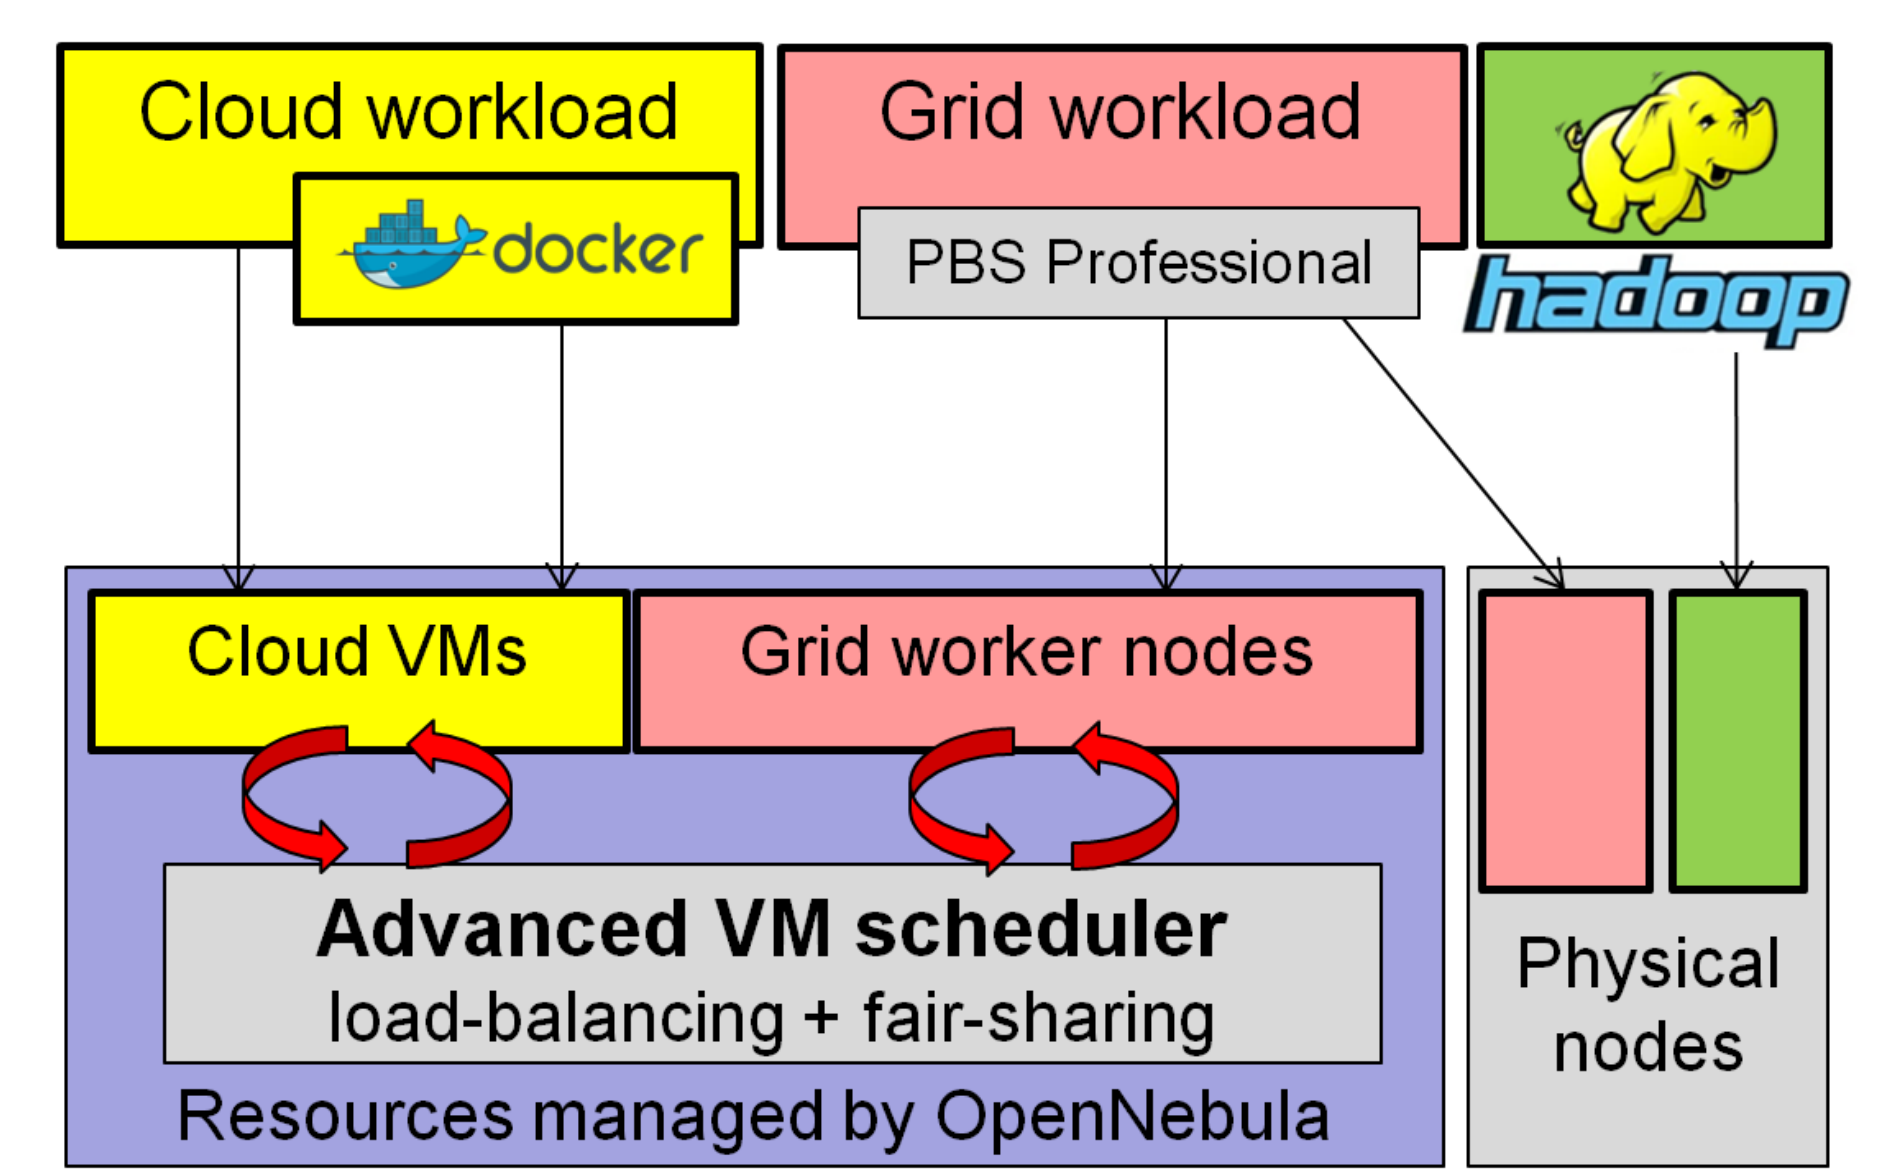
Load-balancing layer:

- global resource manager managing applications’ shares (e.g. MESOS)
- applications must be modified
- yet another framework



Proposed cloud-integrated solution:

- no new resource manager
- load-balancing via cloud layer
- using cloud VM scheduler



5. CONCLUSION AND FUTURE WORK

Current status:

- large portion of the infrastructure is managed by OpenNebula middleware
- hybrid workloads (cloud and grid)
- default VM scheduler is used
- new advanced scheduler is tested [3]

Future work includes:

- automated load-balancing using advanced VM scheduler
- cross-application fair-sharing
- active resource reclaiming
- fair-share driven overbooking

ACKNOWLEDGMENTS

We kindly acknowledge the support provided by the MetaCentrum under the program LM2015042 and the CERIT Scientific Cloud under the program LM2015085. We also highly appreciate the access to MetaCentrum and CERIT Scientific Cloud workload traces.

REFERENCES

- [1] G. Podolníková, B. Parák and D. Klusáček. Extensible and Modular Cloud Scheduler for OpenNebula. In *Cracow Grid Workshop*, 2015.
- [2] D. Klusáček and H. Rudová. Multi-Resource Aware Fairsharing for Heterogeneous Systems. In *Job Scheduling Strategies for Parallel Processing*, 2015.
- [3] G. Podolníková, A. Üрге and D. Klusáček. ONEScheduler: a cloud scheduler for OpenNebula. <https://github.com/CESNET/ONEScheduler>, CESNET, 2016.