MetaCentrum

CESNET

Scheduling Challenges in a Shared Private Cloud Infrastructure

Dalibor Klusáček
Boris Parák
Lukáš Hejtmánek

Cracow Grid Workshop 2017

- **Scheduling of various applications over shared cloud infrastructure in *MetaCentrum***
  - Virtualization enables for easy (de)allocation of resources for
    - **Cloud VMs**
    - **Grid jobs** (running inside “grid worker” VMs)
    - + **Docker** and **Singularity** container engines

application level
(grid & cloud)

static (“by hand”) resource allocation for each application

infrastructure level

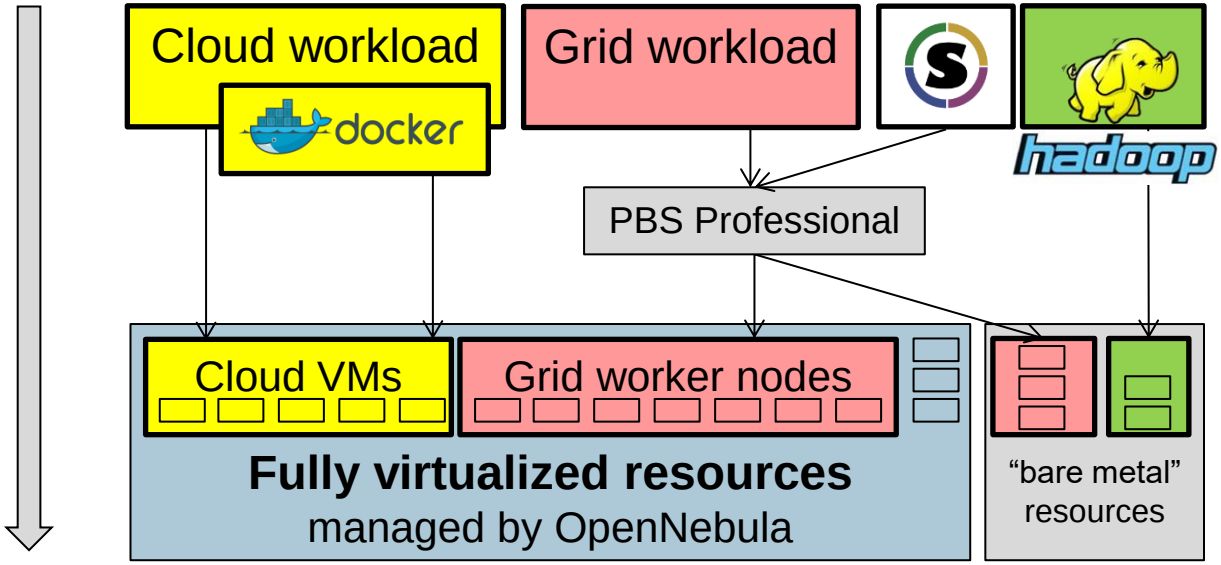

Scheduling Challenges

- **Applications' demands grow in time**
- **"Free of charge" computing** for scientists & academia
- **Cloud computing then involves many challenges**
  - **Resource reclaiming** (it's free => many "zombie VMs")
  - **Fairness** (classic fairshare-like mechanisms are not suitable)
  - **Resource utilization** (e.g., CPU load) in the cloud is not high

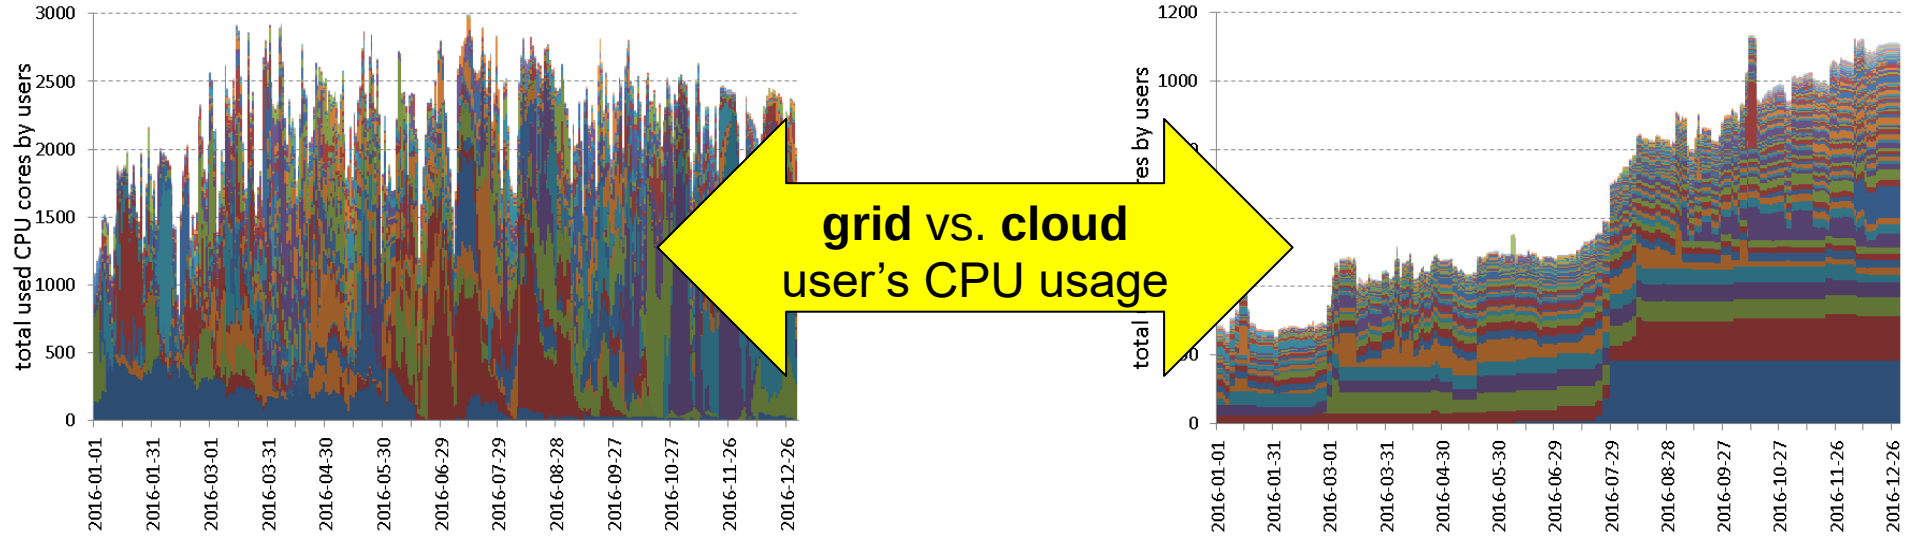

Applied and Considered Solutions

- **More aggressive overbooking** based on observed data
  - Improved CPU utilization (~15% → ~50%)
  - Still leaving space for “elasticity”
- **Default VM lifetime set to 3 months**
  - With the possibility of user-initiated lifetime extension
  - Eliminated ~20% of old VMs
- **Considered improvements**
  - **“Scavenge-like” grid computing** on idle cloud nodes
  - **User and VM prioritization/classification** to reflect fairness and various SLOs
  - **Dynamic VM overbooking** based on given criteria (fairness)
  - In general, **advanced scheduling techniques** used to orchestrate overbooking, live migrations (rescheduling) and VM prioritization