

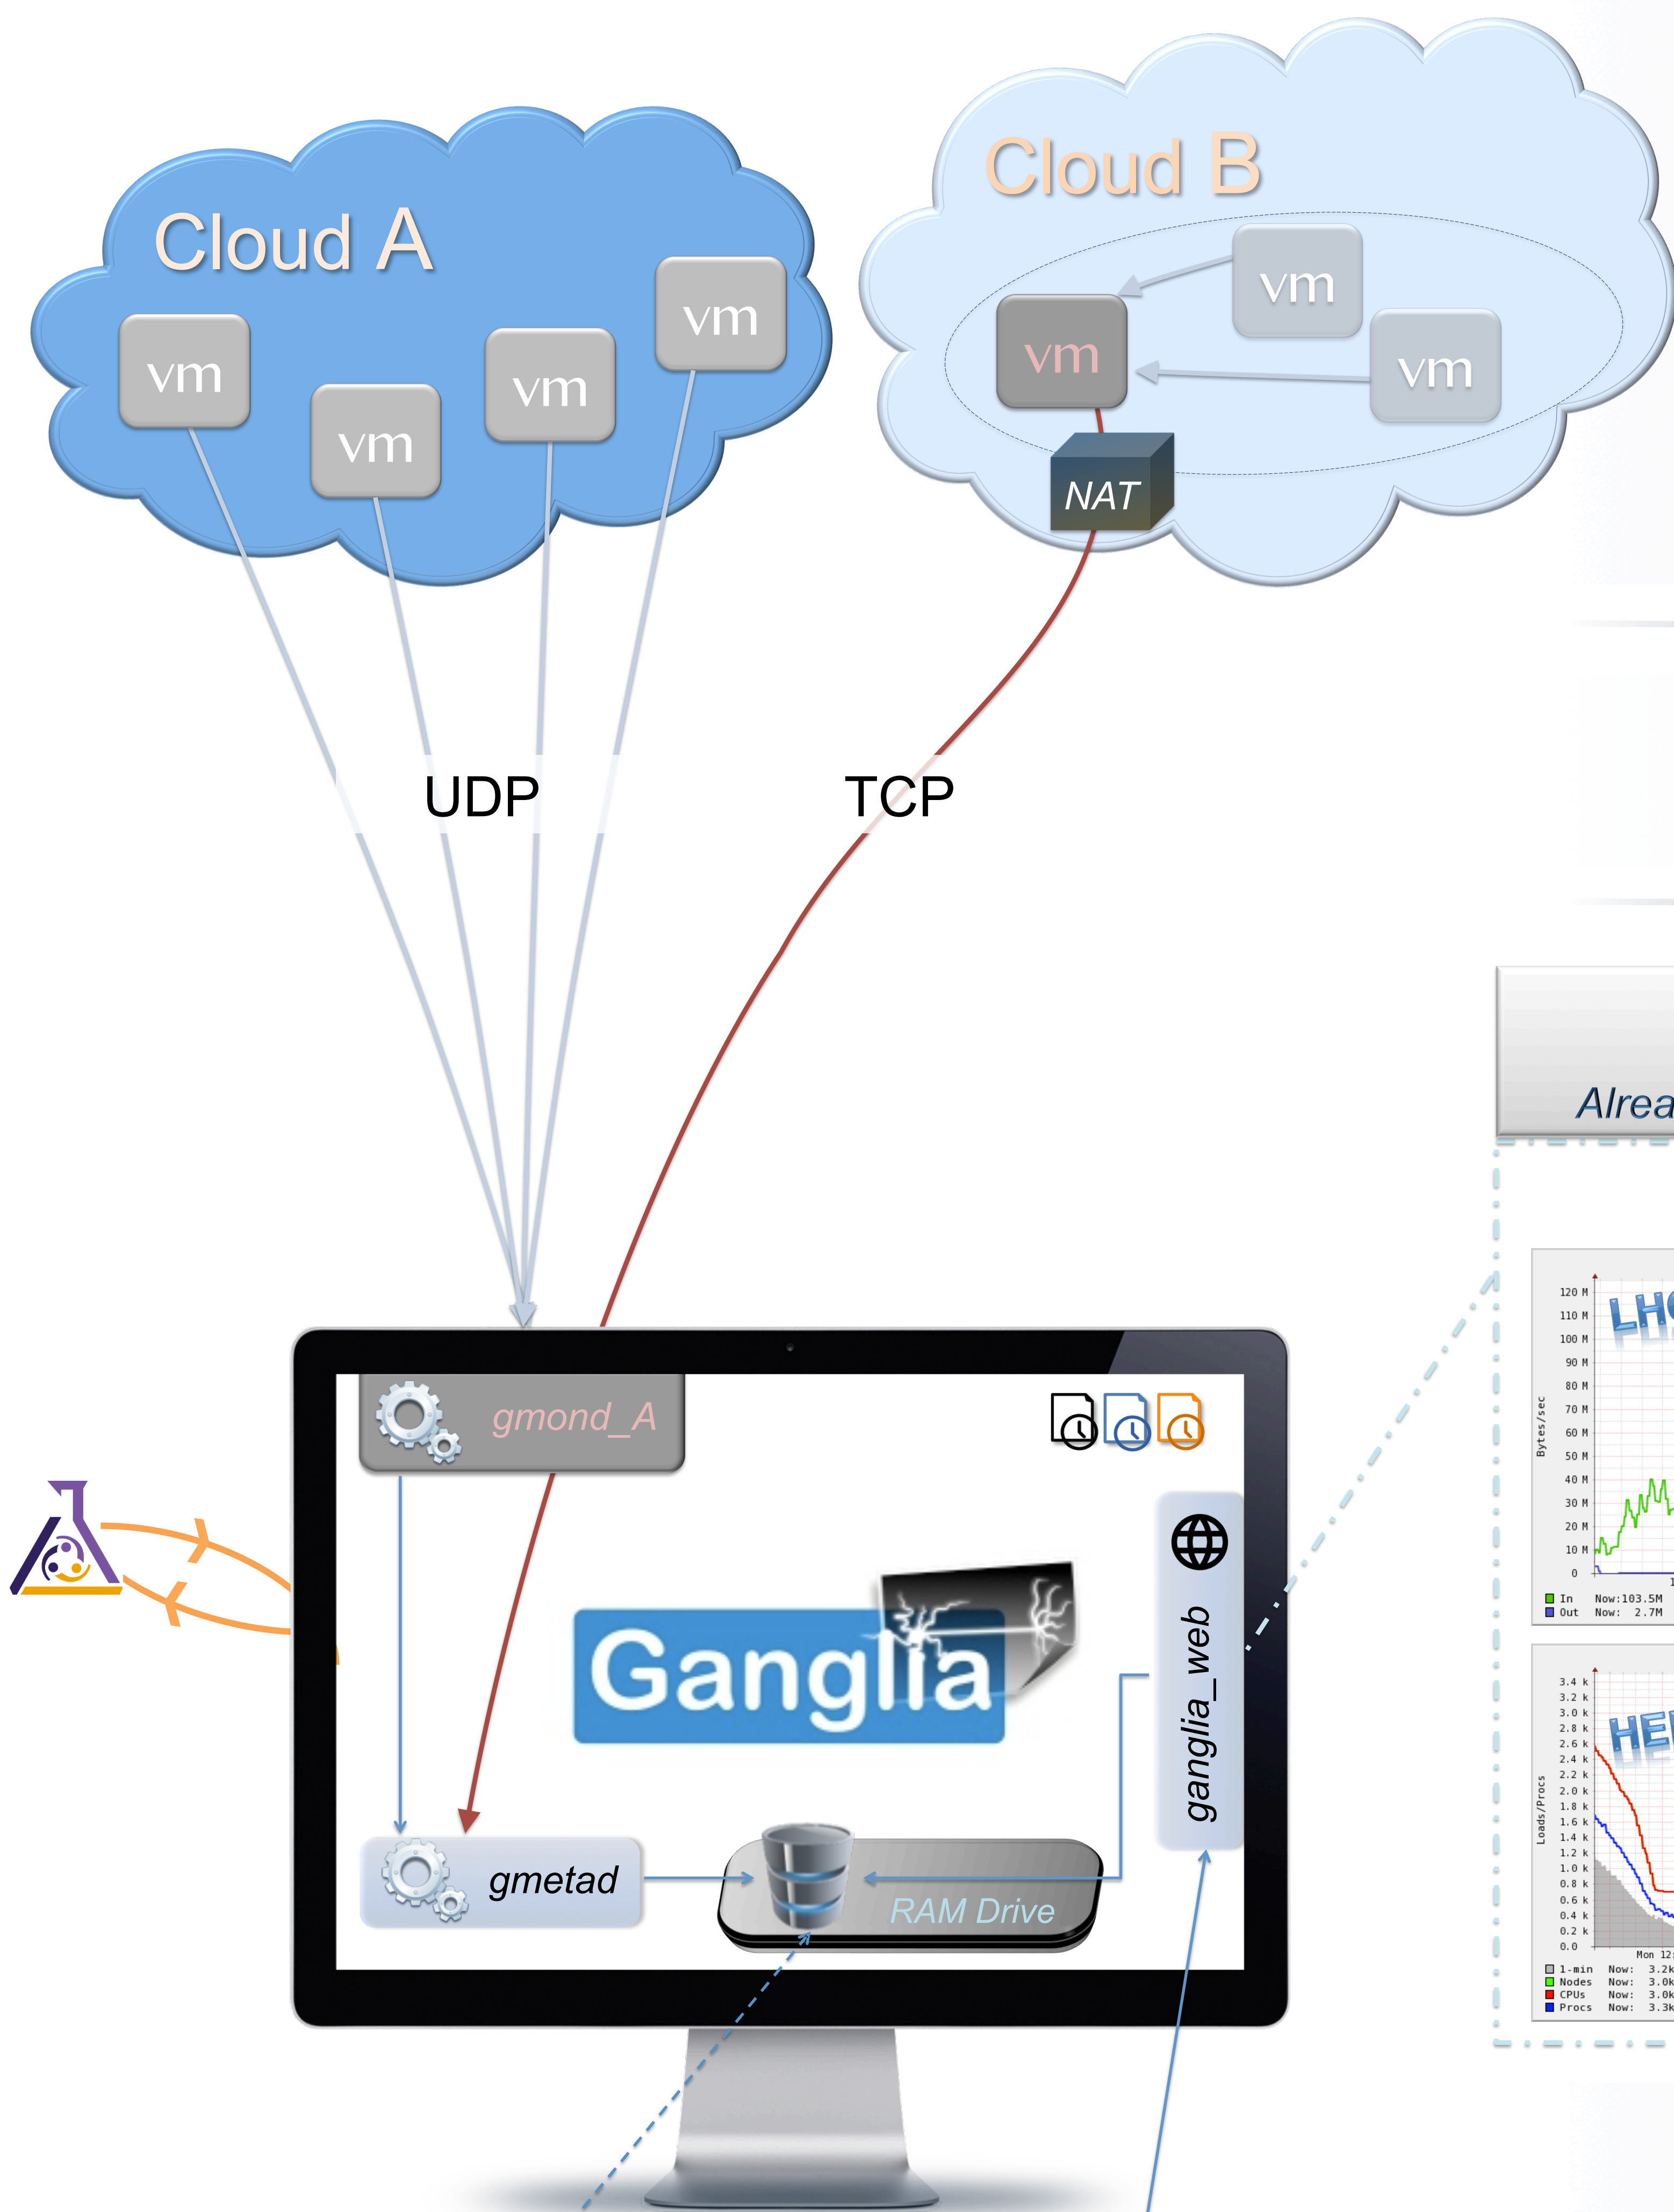
Monitoring the Delivery of Virtualized Resources to the LHC Experiments

C. Cordeiro¹, A. Di Girolamo¹, L. Field¹, D. Giordano¹, D. Spiga¹, L. Villazon^{1,2}

¹CERN, ²Universidad de Oviedo



The adoption of Cloud technologies by the LHC experiments puts the burden of monitoring virtual resources upon the VO. **Monitoring the instantiated virtual machines** is therefore a **fundamental activity** and here it is described how the **Ganglia monitoring system**^[1] can be exploited for monitoring and also for **providing reliable information for other applications**.

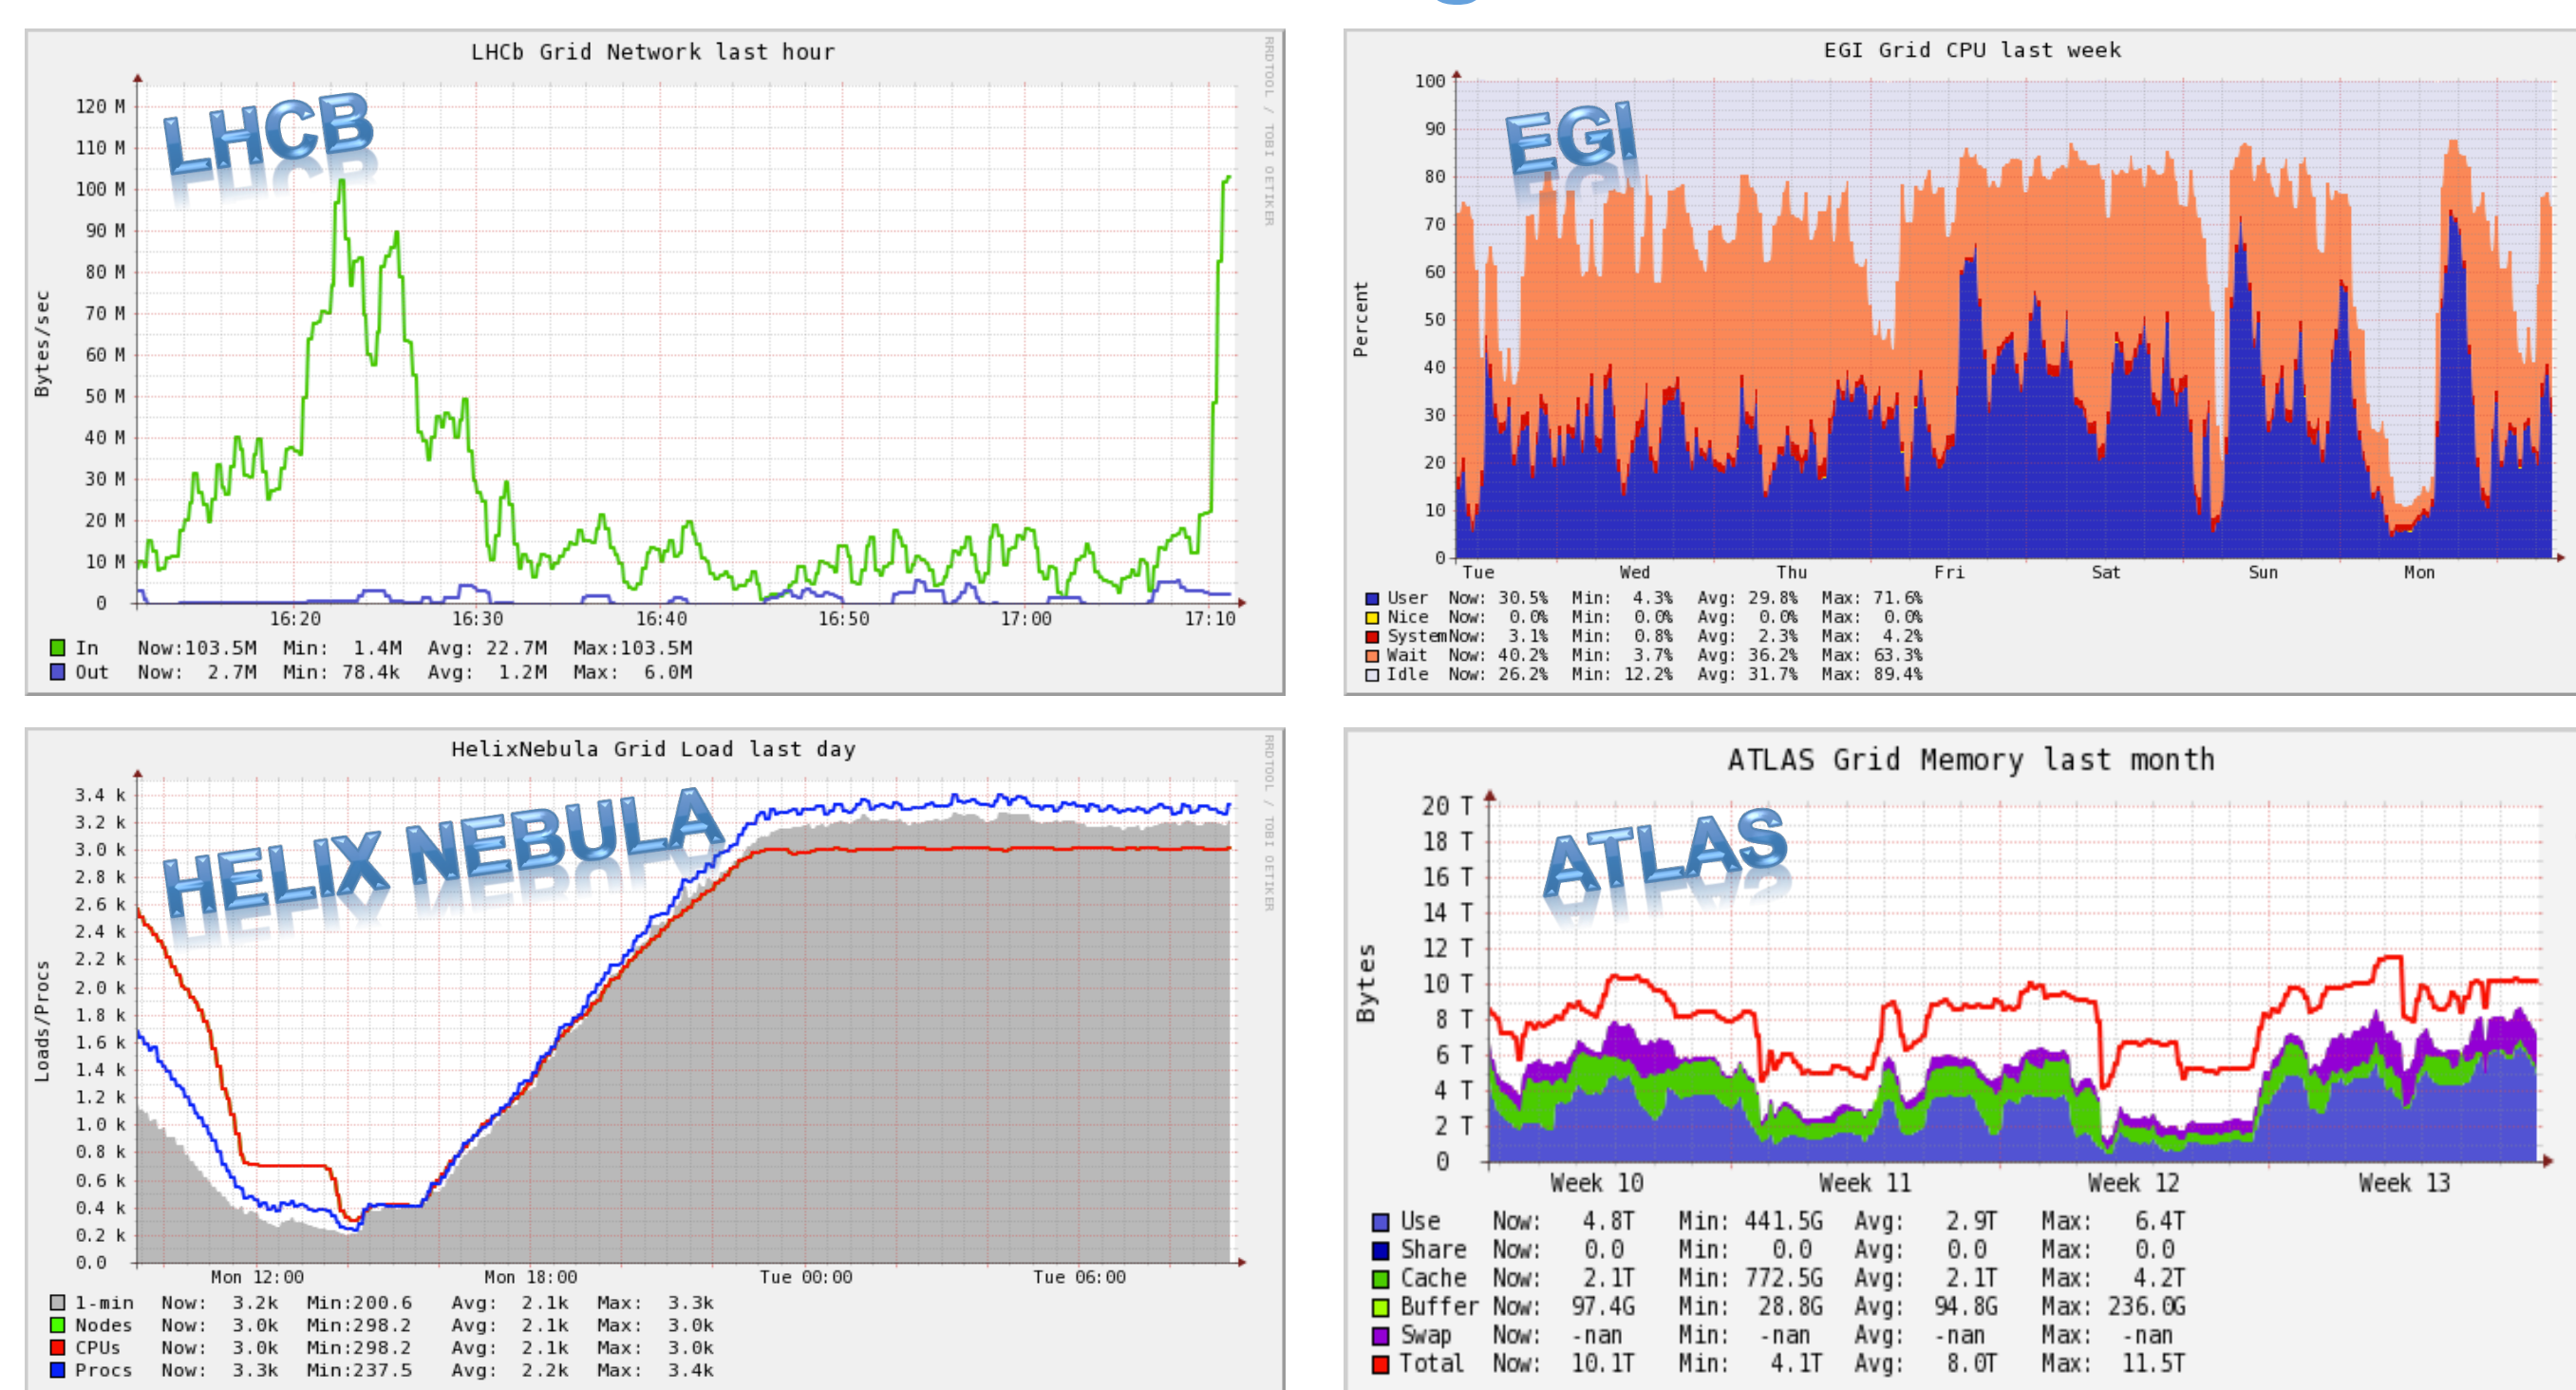


- Ganglia's monitoring daemon (*gmond*) sits on every Virtual Machine (vm), gathering monitoring statistics and sending them to a receiver node (vm) through unicast
- Each provider has its own *gmond* configuration, which relies on a **port number**, **cluster name** and **host** (collecting *gmond*) **address**
- To configure *gmond*, one can do it dynamically by fetching a pre-built JSON file, with *gmond*'s configuration from the *gmetad*'s server during the VM's contextualization

- Communications between *gmond* are done via UDP
- From *gmond* to *gmetad* these are done via TCP

Deployed for **ATLAS**, **LHCb**, **CMS**^[2] and **Helix Nebula**^[3]
Monitoring **>15** cloud providers
Already **extracted** and **stored ~440 GB** of raw metrics (not aggregated)

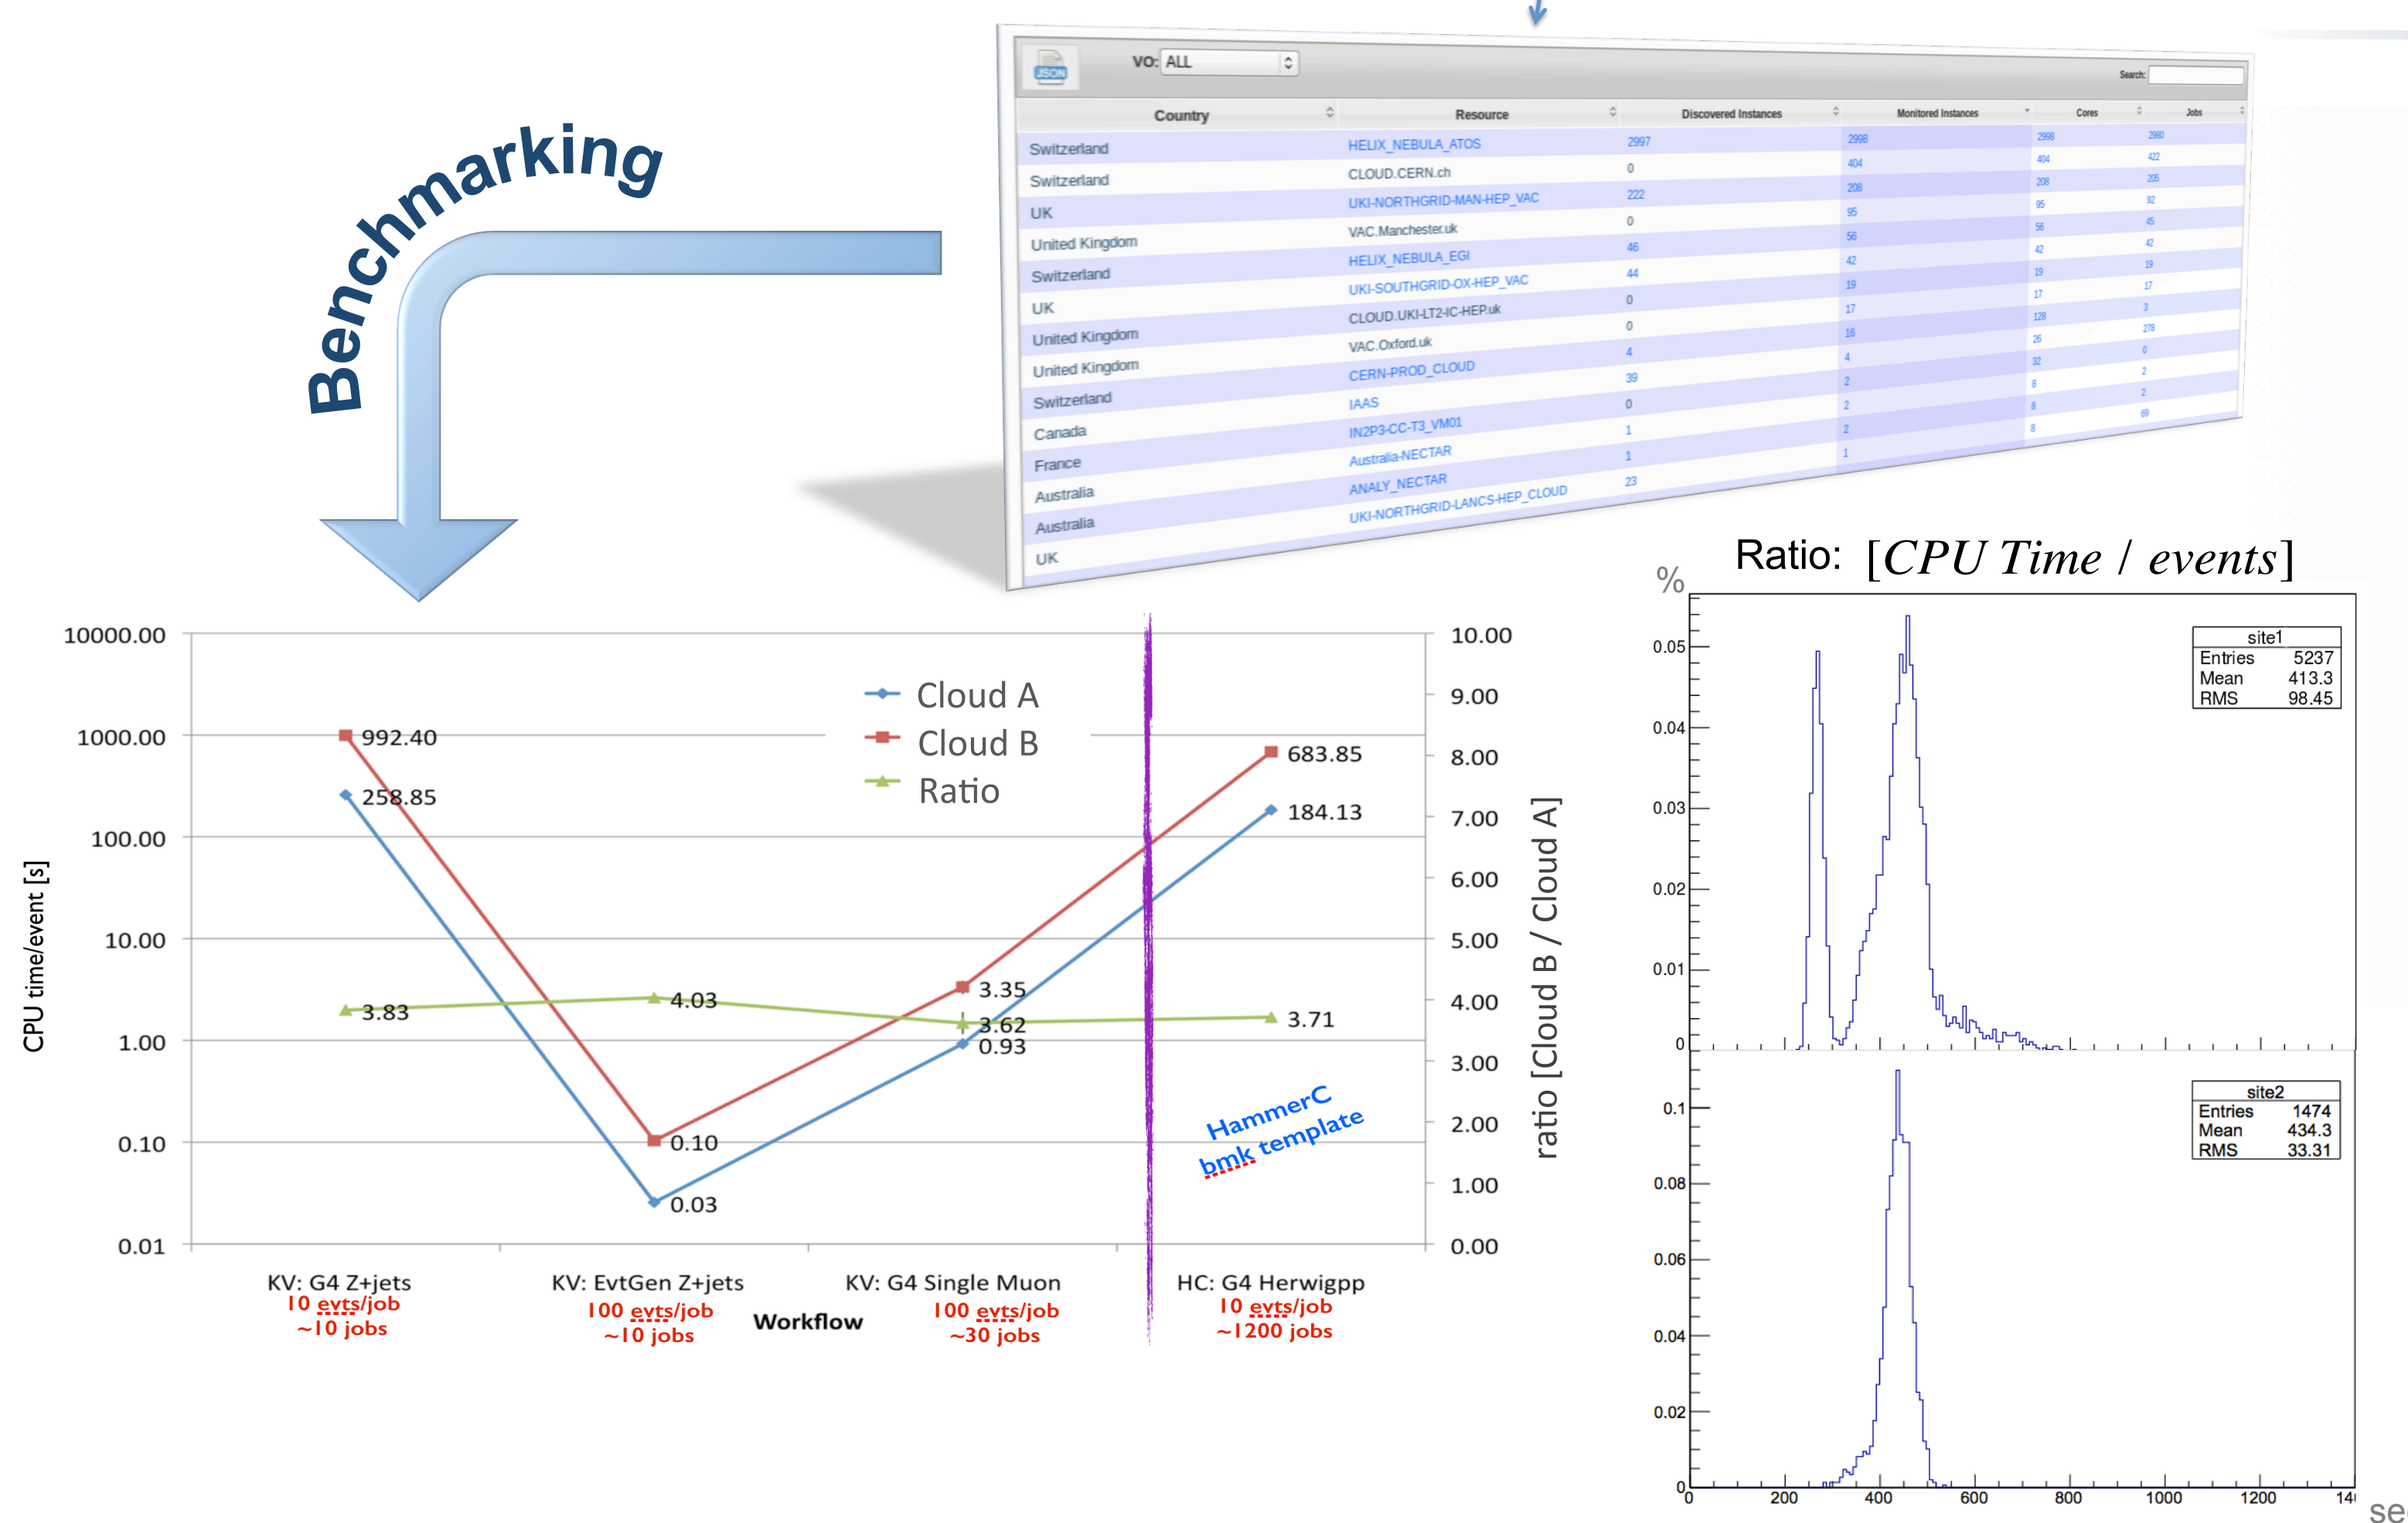
Cloud monitoring overview



alarm RAM Drive usage : clean dead hosts from RRD DB : backup the RRD DB

- The Ganglia server is deployed, configured and managed through Puppet
- Ganglia's *gmetad* daemon polls *gmonds* periodically and stores the metrics in a RRD database sitting inside a dedicated RAM drive
- Ganglia Web accesses the RRD DB and generates the web UI
- Time-series monitoring data can afterwards be extracted and repurposed to serve consumer accounting, data analytics and others

Benchmarking



- Ganglia's data is used by an Accounting system^[4] as a reliable source for measuring the resource usage from the consumer's perspective
- Later, the comparison of consumer/provider and cloud/site accounting triggers the need for resource profiling
- Since most times (within clouds) resources are seen as black boxes, benchmarking computing resources w.r.t. experiment workflows becomes essential in order to evaluate performance and costs
- With the undergoing adoption of clouds at the WLCG, resource profiling takes an important role in the task of comparing cloud technology with the current GRID servers

References

- Ganglia Monitoring System: <http://ganglia.sourceforge.net/>
- ATLAS, LHCb and CMS Ganglia monitors: <http://agm.cern.ch/>, <http://lqm.cern.ch/>, <http://cgm.cern.ch/>
- "Accessing Commercial Cloud Resources within the European Helix Nebula Cloud Marketplace", CHEP2015 Poster 216
- Cloud Accounting portal: <http://cloud-acc-dev.cern.ch/>

