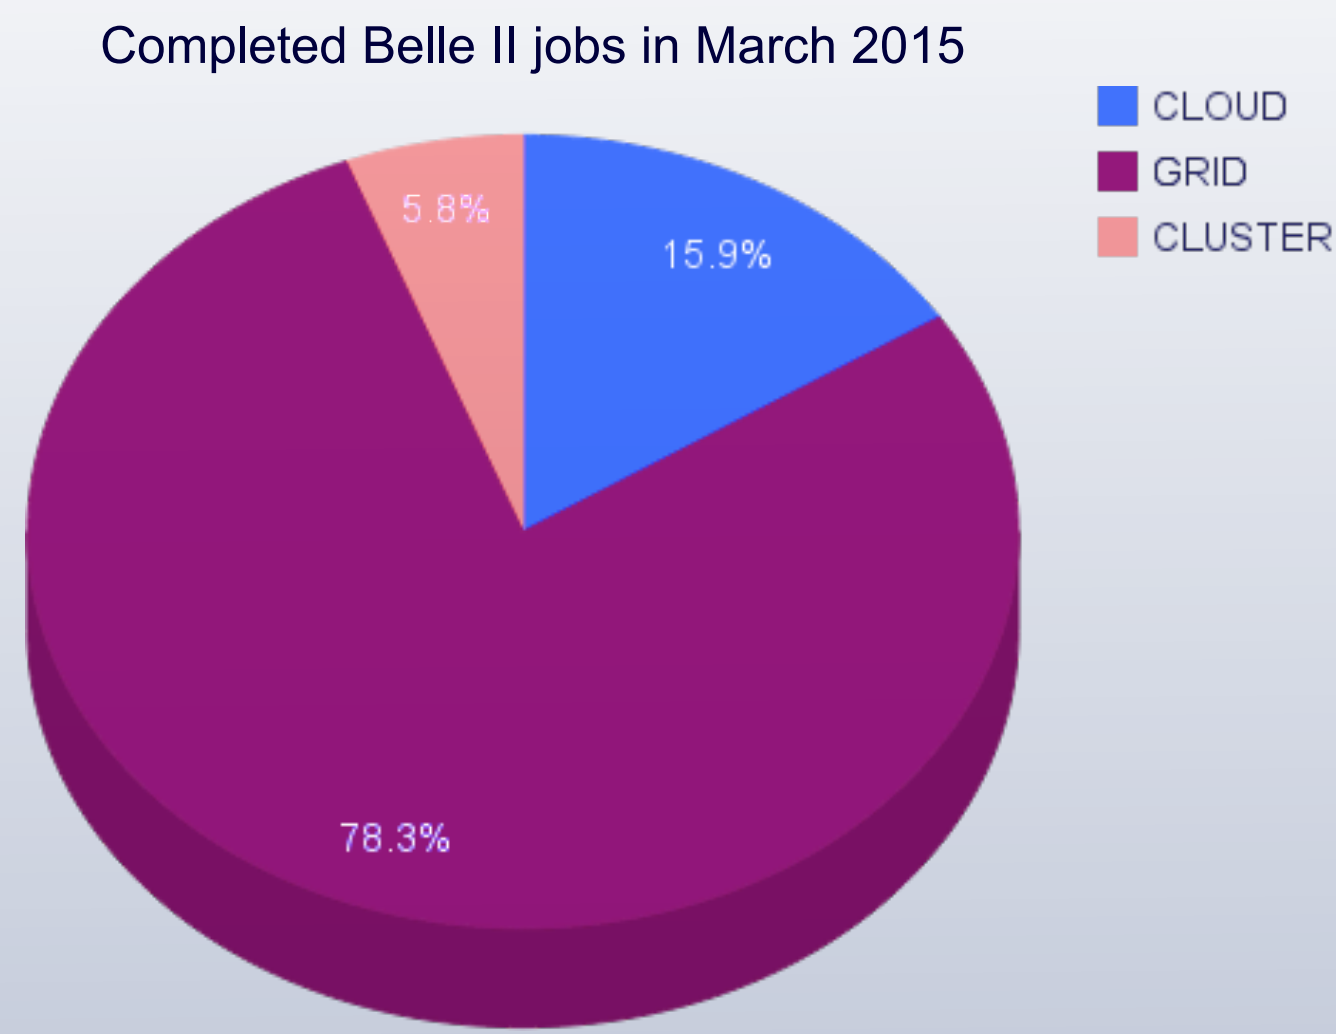
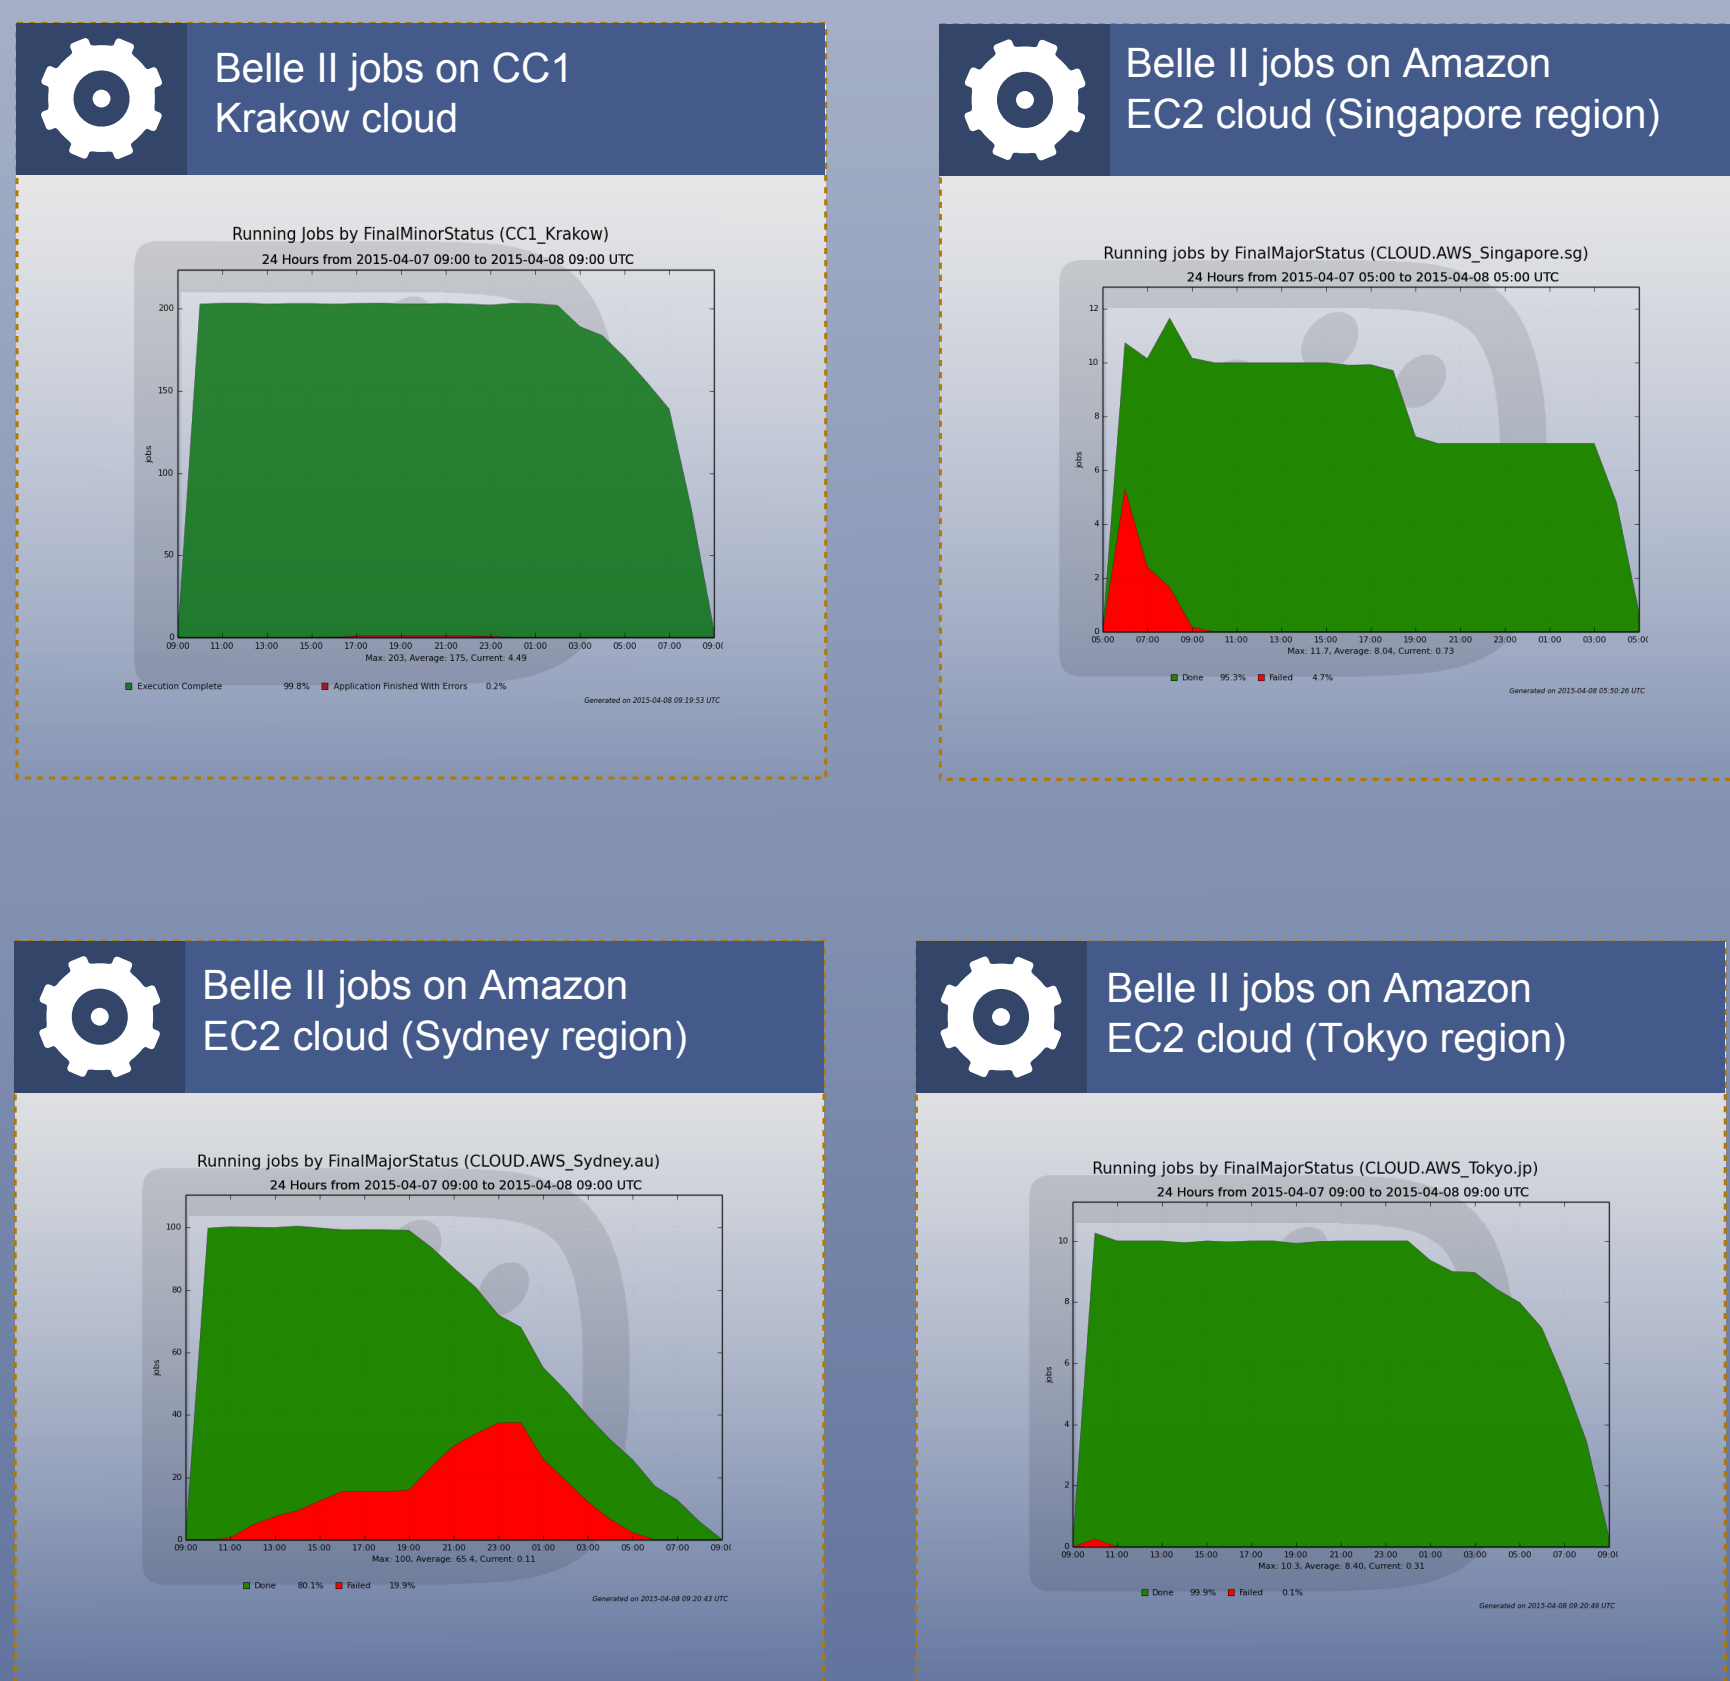




Computing for the Belle II experiment



Belle II experiment is able to compute on variety kinds of resources, i.e. grids, standalone clusters, academic clouds and commercial clouds. Clouds are becoming increasingly important for Belle II computing model^[1]. For now we can handle communication to EC2, OCC1 and Nova APIs. It means that we can run virtual machines on Amazon, CC1, OpenStack and other compatible clouds.



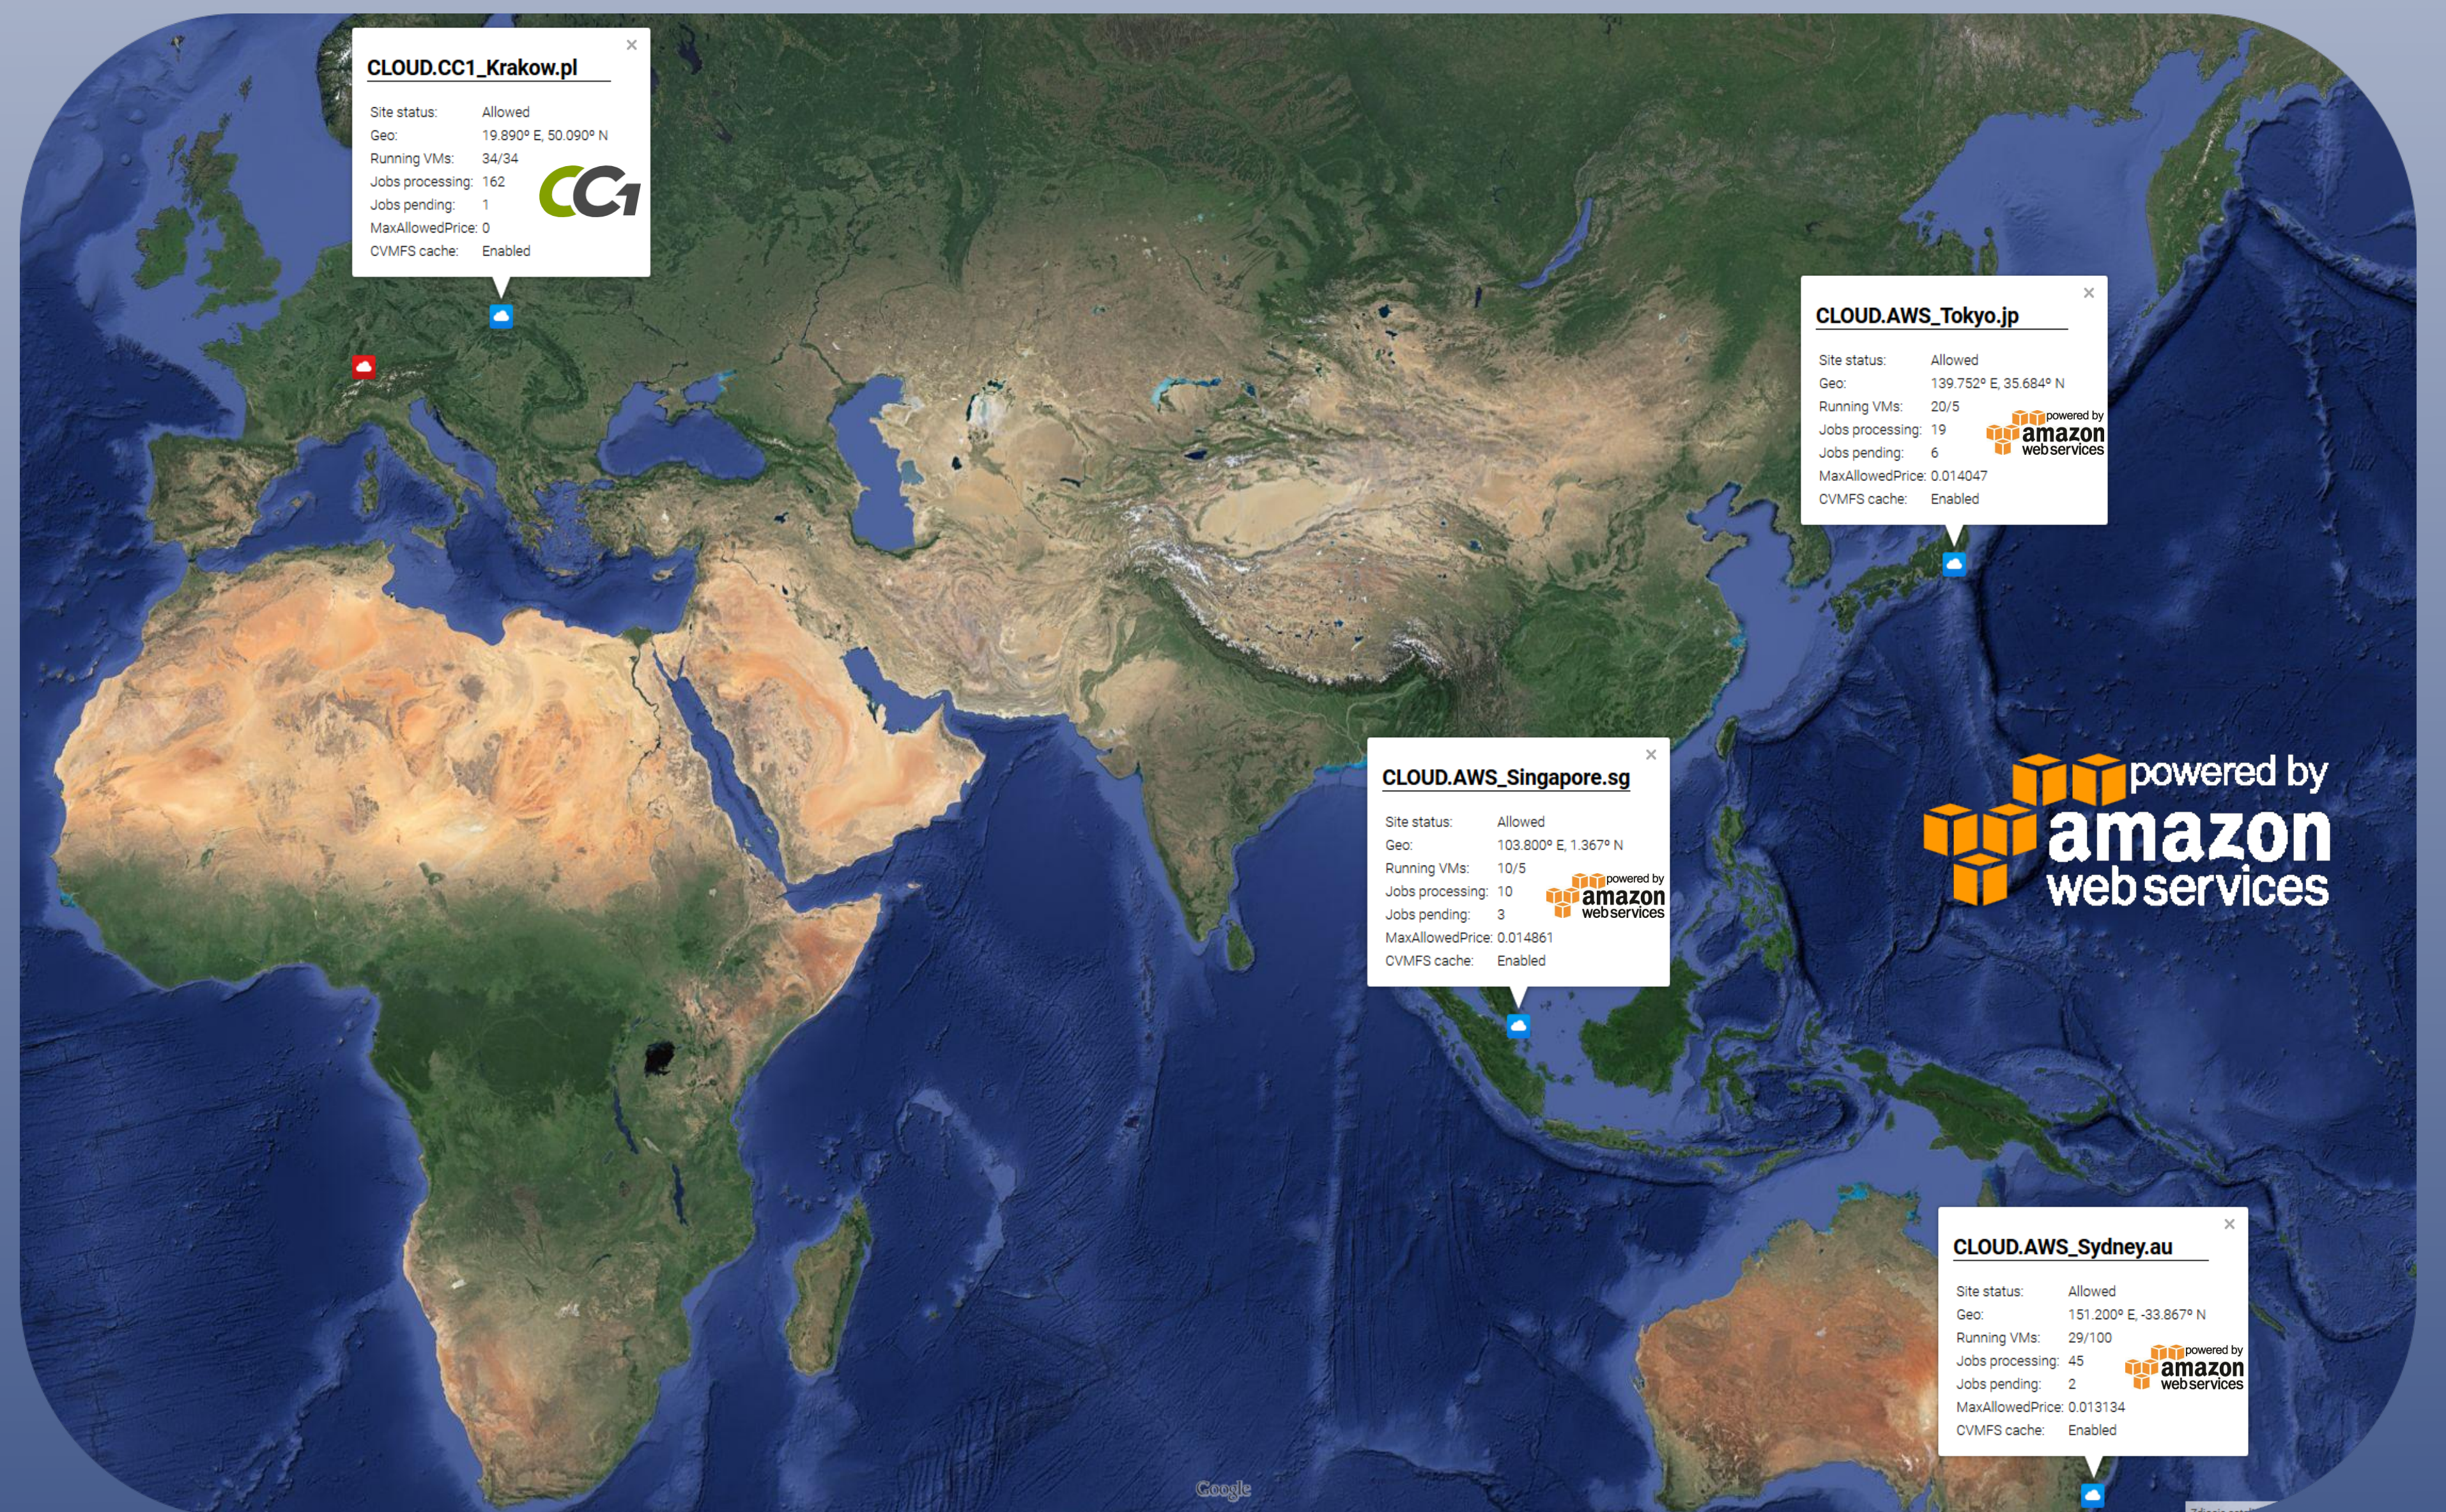
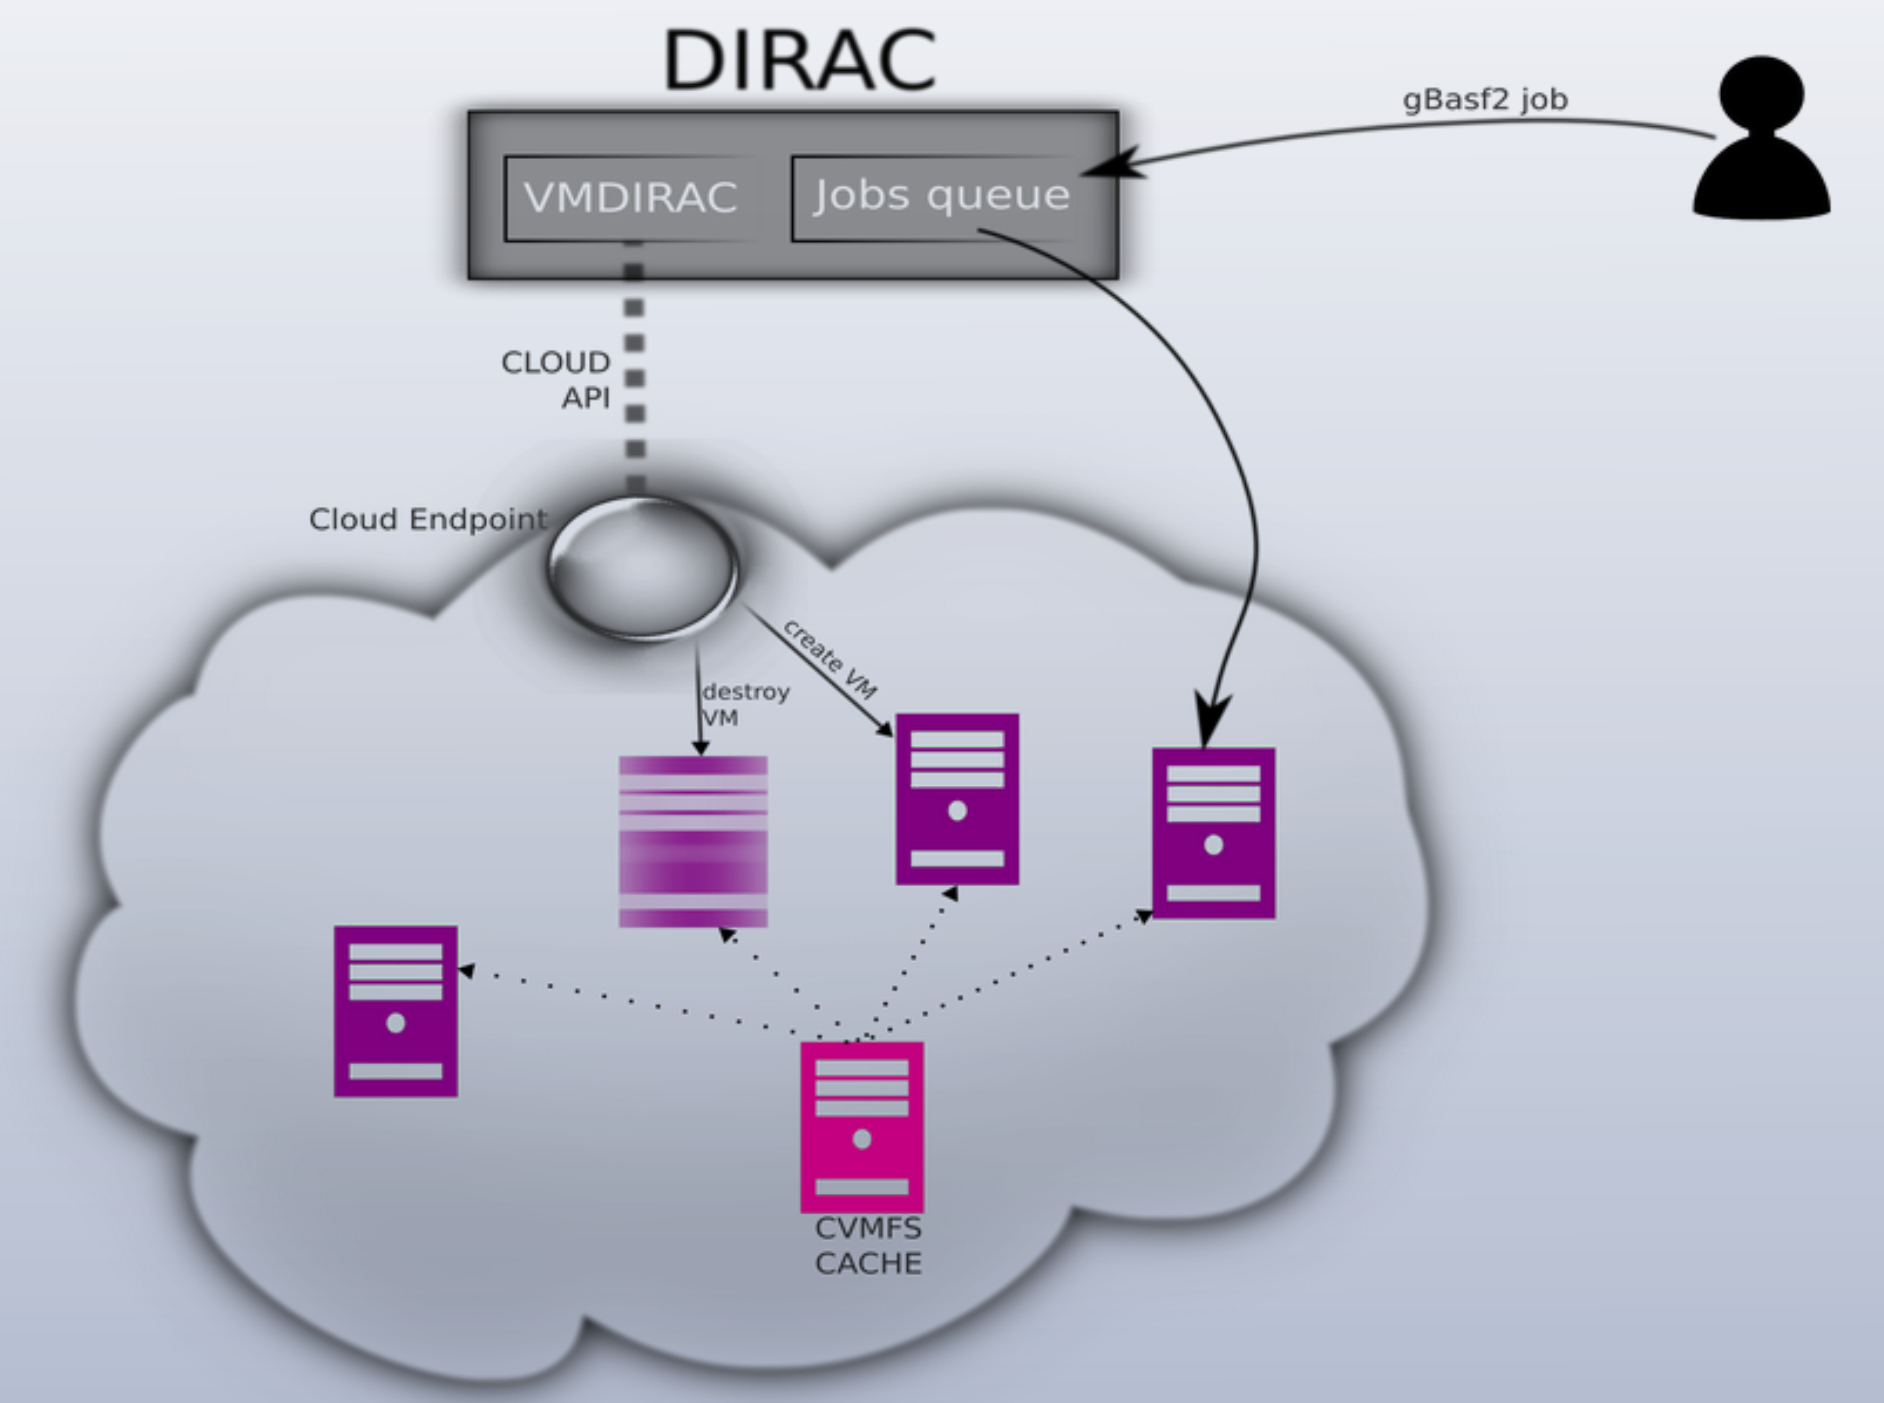
Belle II jobs (gBasf2) can take many hours of CPU so interruptions are crucial. With the DIRAC system we achieved great failover service and even sudden Spot Instances shutdown can be easily handle and job can be rescheduled. Such situation can be observed in one of charts above. AWS Sydney region became expensive in working day time so lots of VMs based on Spot Instances has been destroyed.

VMDIRAC^[2] - Cloud resources manager for the DIRAC system

The DIRAC^[3] (Distributed Infrastructure with Remote Agent Control) project is a complete Grid, Cloud, Host and Volunteer solution for a community of users such as the LHCb Collaboration, Belle II Collaboration or NGI multi-VO portals (DIRAC4EGI, FranceGrilles...).

Cloud resources can be connected by VMDIRAC module through public interfaces. In particular the mechanism of dynamic activation of new virtual machines with reserved job slots for new tasks in case of an increasing demand for computing resources.

VMDIRAC is an extension of DIRAC system to manage pools of different kinds of clouds by drivers to EC2, OCC1, rOCCI, Nova,



Additional Agents for VMDIRAC EC2 driver

EC2UserDataAgent

EC2 machines are contextualized by user_data script. user_data script for cloud-init is generated periodically from the DIRAC Configuration.

EC2SpotPriceAgent

This agent is used to calculate optimal MaxBiddingPrice for Spot Instances requests in each Amazon region and zone. Also increase VMs limits for the cheapest zone and decrease VMs limits for expensive zones. Results are saved in the main DIRAC Configuration.

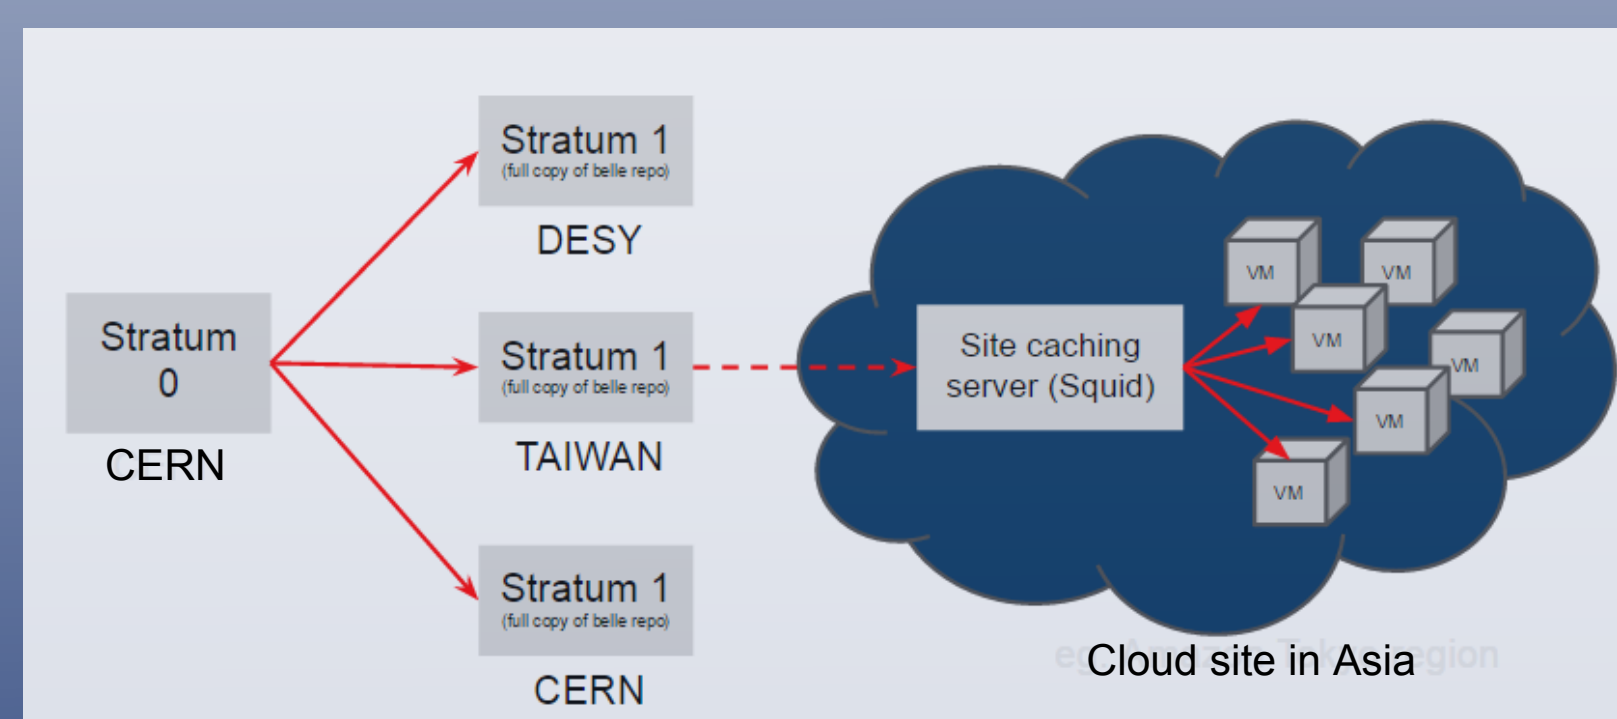


CVMFS^[5] Belle II software repository

Belle II software is located on the same mirror servers which supports LHC experiments this give us more stable and up-to-date access to experiment's software.

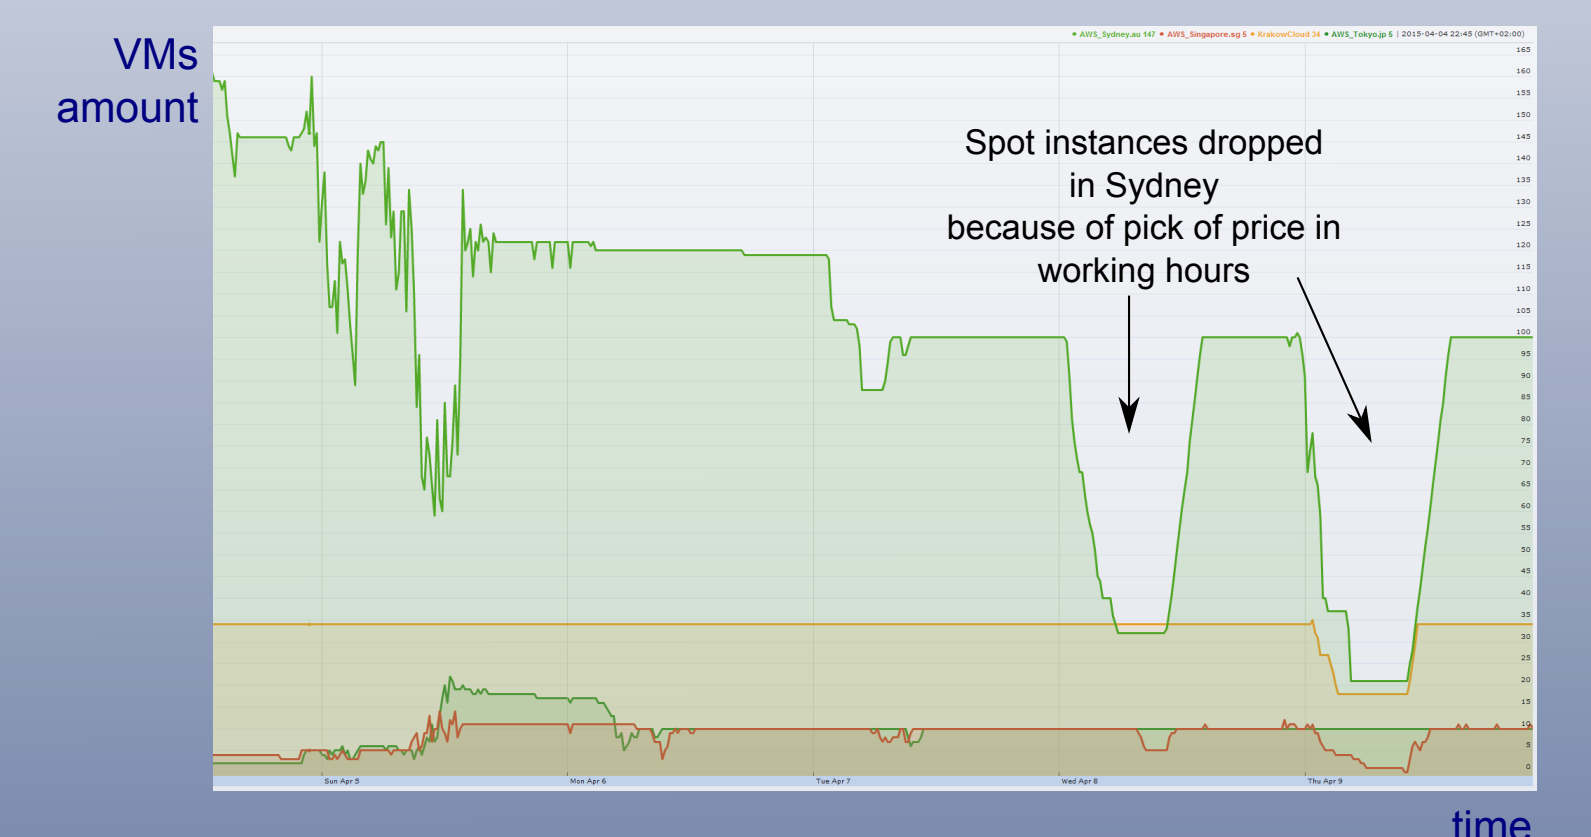
Squid proxy server on each cloud region (site) managed by Belle's VMDIRAC is set up for caching software mounted by CVMFS to virtual machines.

Amazon On-demand (persistent) instance type is used for Squid servers - not Spot instances which can suddenly disappear.



Summary

- Implemented costs optimization engine for the Amazon EC2 public cloud works well.
- Active VMDIRAC/BelleDIRAC development to provide environment for cloud processing, management and monitoring.
- Development installation at Institute of Nuclear Physics PAN in Krakow.



References

- [1] Belle, "Belle II Technical Design Report", arXiv:1011.0352, (2010)
- [2] "Cloud flexibility using DIRAC interware", Victor Fernandez Albor et al 2014 J. Phys.: Conf. Ser. 513 032031
- [3] DIRAC INTERWARE, <http://diracgrid.org>
- [4] "Powering Distributed Applications with DIRAC Engine", The International Symposium on Grids and Clouds (ISGC) 2014, March 23-28, 2014, Academia Sinica, Taipei, Taiwan
- [5] CVMFS, <http://cernvm.cern.ch/portal/filesystem>