UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Secure Hybrid Cloud Infrastructure for Scientific Applications

Project Members:

Paula Eerola
Miika Komu
Matti Kortelainen
Tomas Lindén
Lirim Osmani
Sasu Tarkoma
Salman Toor (Presenter) *salman.toor@helsinki.fi*
John White

DII-HEP project

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Datacenter Indirection Infrastructure for Secure HEP Data Analysis

- Collaboration between
  - Helsinki Institute of Physics (HIP)
  - Computer Science Department, University of Helsinki
- Funded by Academy of Finland 2012 - 2014

Goals of the project

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

- Explore the latest software stacks for distributed computing infrastructures
- Construct a secure and scalable setup for scientific applications
- Use the CMS analysis and production framework as a test case

Initial setup

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

- Jade 768 cores main ARC CE, shared by CMS and ALICE
  - Madhatter dCache Storage Element connected to Jade
- Alcyone 892 cores ARC CE, shared resource, 10 km from Jade
- Korundi 400 cores ARC CE, shared resource, 10 km from Jade

Cloud based setup

- We have a production site based on private cloud setup
  - OpenStack (Folsom → Icehouse)
  - Gluster Filesystem
  - Advanced Resource Connector (ARC) middleware for providing grid interfaces
  - CERN VM File System (CVMFS)
  - OpenStack deployed on Ubuntu 12.04 LTS
  - VMs based on Scientic Linux CERN 6.4

System architecture

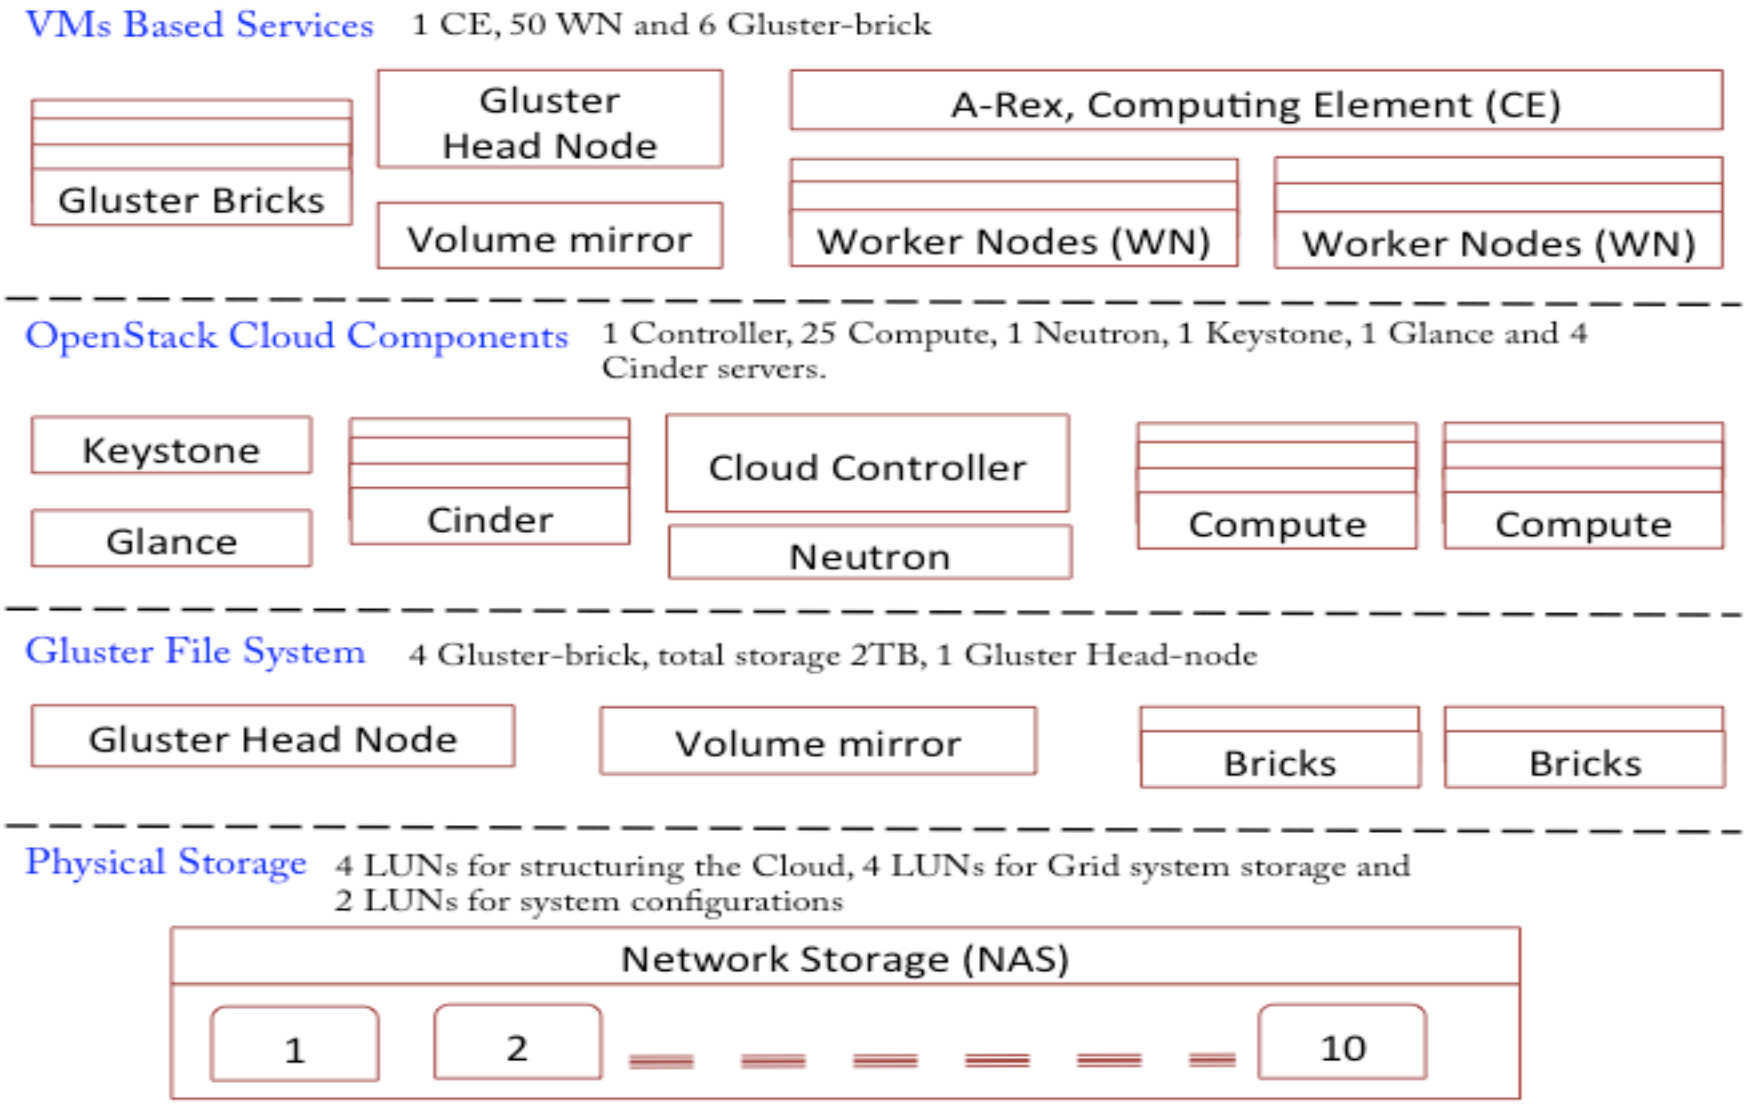

Site status

- So far the system has run a maximum of 200 concurrent jobs
- The average CPU efficiency of all CMS jobs over the previous year is *83%* from the Swedish Grid Accounting System (SGAS).

Machine view for nodeslab-0002.nlab.tb.hiit.fi

Start month 2013-08 End month 2014-08 Submit

Top 10 projects for the selected date range

	Walltime days	Efficiency	Number of jobs
cms	9163	83	147290

Performance analysis

- Application Level
  - 4% performance loss evaluated with the HEPSPEC-2006 (Thanks to Ulf Tigerstedt, CSC for help with HEPSPEC tests)
- System Level
  - VM boot response both at local vs GlusterFS based setup

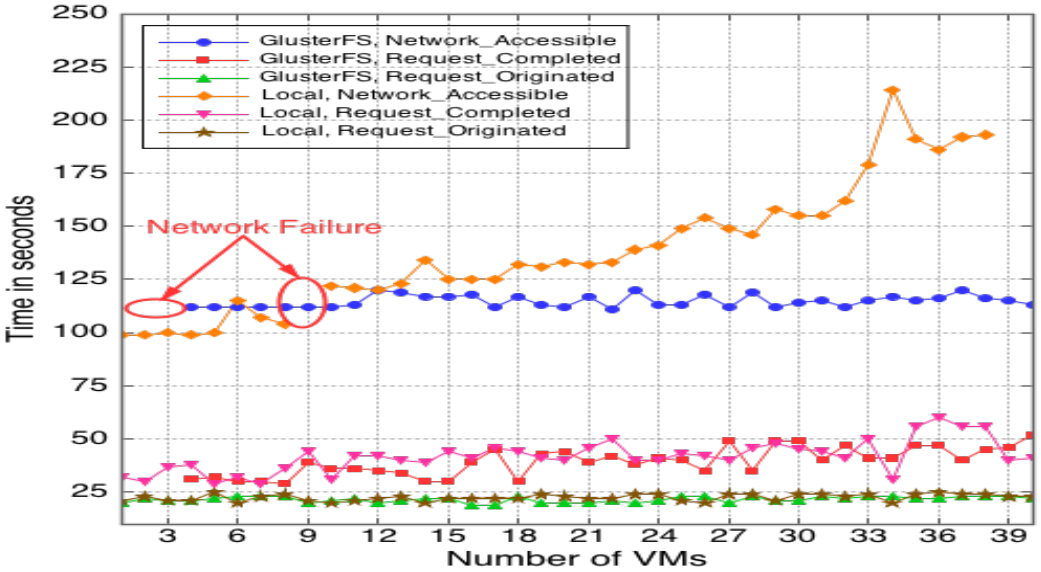

GlusterFS IO response

- Brick 1 and 2 is used inside the Cloud
- Brick 3 and 4 is used for the Glance and Nova

	Brick-1	**Brick-2**	**Brick-3**	**Brick-4**
Days	20	20	20	20
Total Reads	40GB	41GB	169GB	831GB
Total Writes	38GB	36GB	364GB	1017GB
Average – Maximum Latency (milliseconds) / No. of Calls (in millions)				
Read	0.05-11.6/ 29.8	0.05-10.8/ 31.0	0.29-217.1/ 0.6	0.38-2386/ 0.32
Write	0.08-16.8/ 1.8	0.07-14.0/ 1.9	1.66-5151/ 10.8	1.06-7281/ 7.8

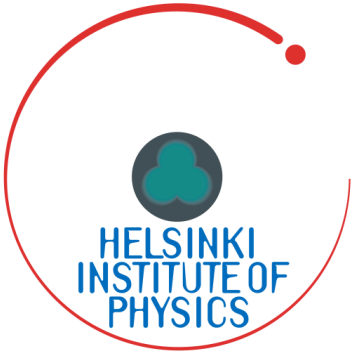

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Article in CHEP-2013:

S. Toor, L. Osmani, P. Eerola, O. Kraemer, T. Lindén, S. Tarkoma, J. White. **A scalable infrastructure for CMS data analysis based on OpenStack Cloud and Gluster file system**. Journal of Physics: Conference Series 513 062047 doi:10.1088/1742-6596/513/6/062047.

HELSINKI INSTITUTE OF PHYSICS

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Elastic solution

- We have extended the EES plugin of the EMI Argus service
- Execution Environment Service (EES) originally designed for the OpenNebula but now it can work with OpenStack Cloud as well
- *John White* is mainly responsible for this work
- Argus client-side modifications completed
- The component is not production ready yet

OpenStack, Argus and EES extension

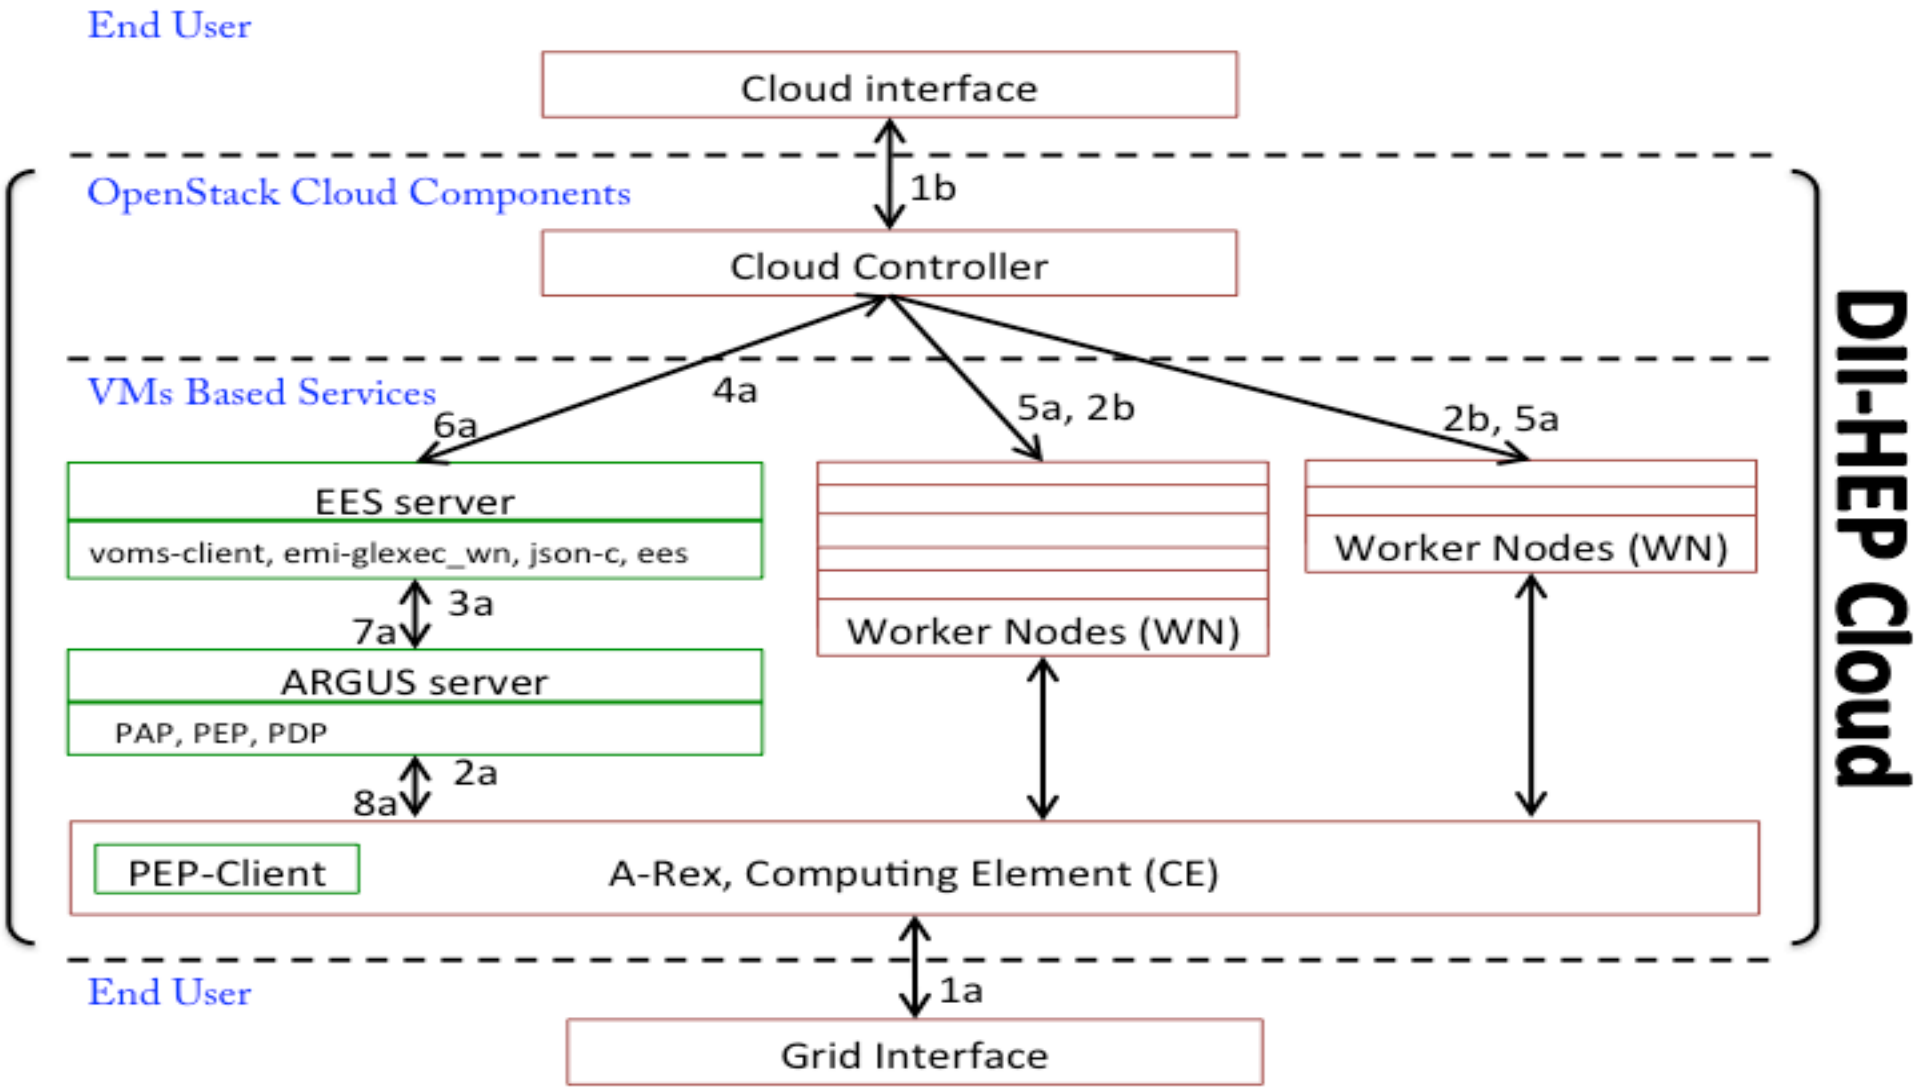

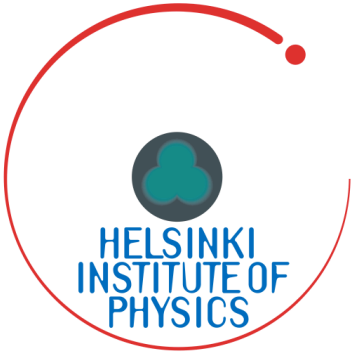

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Article in ISGC-2014:

J. White, S. Toor, L. Osmani, P. Eerola, O. Kraemer, T. Lindén, and S. Tarkoma**. Dynamic Provisioning of Resource in a Hybrid Infrastructure**. In press, International Symposium on Grids and Clouds (ISGC) 2014, Taipei, Taiwan.

Secure cloud setup

- We have used Host Identity Protocol (HIP) for structuring the secure cloud

- Host Identity Protocol (HIP)
  - Designed for mobile networks
  - Provides a secure mechanism for IP multihoming and mobility (VM migration)
  - HIP separates the end-point identifier and locator roles of IP addresses
  - Provides persistent cryptographic identifiers
  - Supports both IPv4 and IPv6 addressing
  - The Host Identifiers (HI) are not routable, so they are translated into routable addresses (locators) between network and transport layer
  - The HIP connections are typically protected with IPSec

  HIP -> http://infrahip.hiit.fi

Secure private cloud

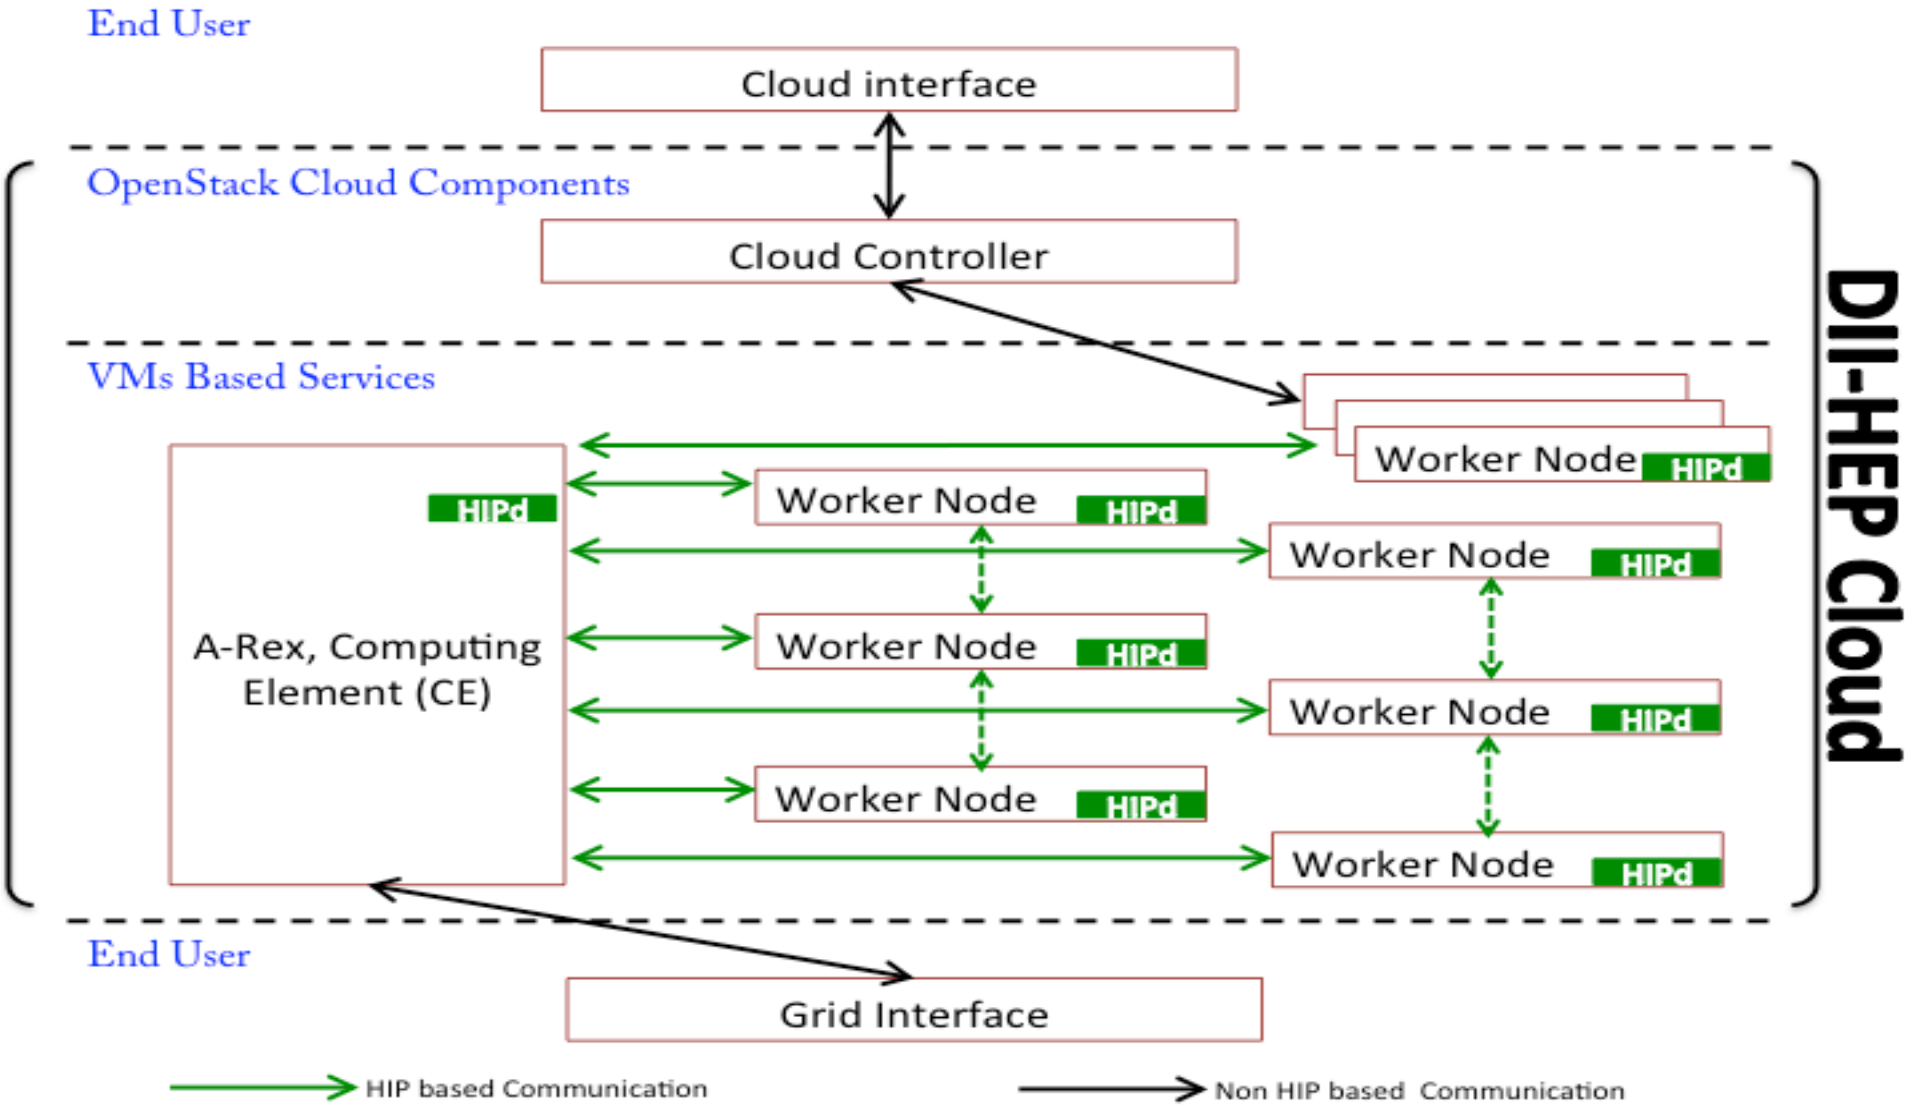

Secure hybrid cloud

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

- CSC, The Finnish IT Center for Science has a new data center in Kajaani, more than 500 km from Helsinki
- CSC has a HP 9216 cores (Q4 17000 cores) supercluster Taito, with 16 cores and 64 GB RAM on most nodes
- The Pouta cloud is an Infrastructure as a Service running on a part of Taito

Secure hybrid cloud

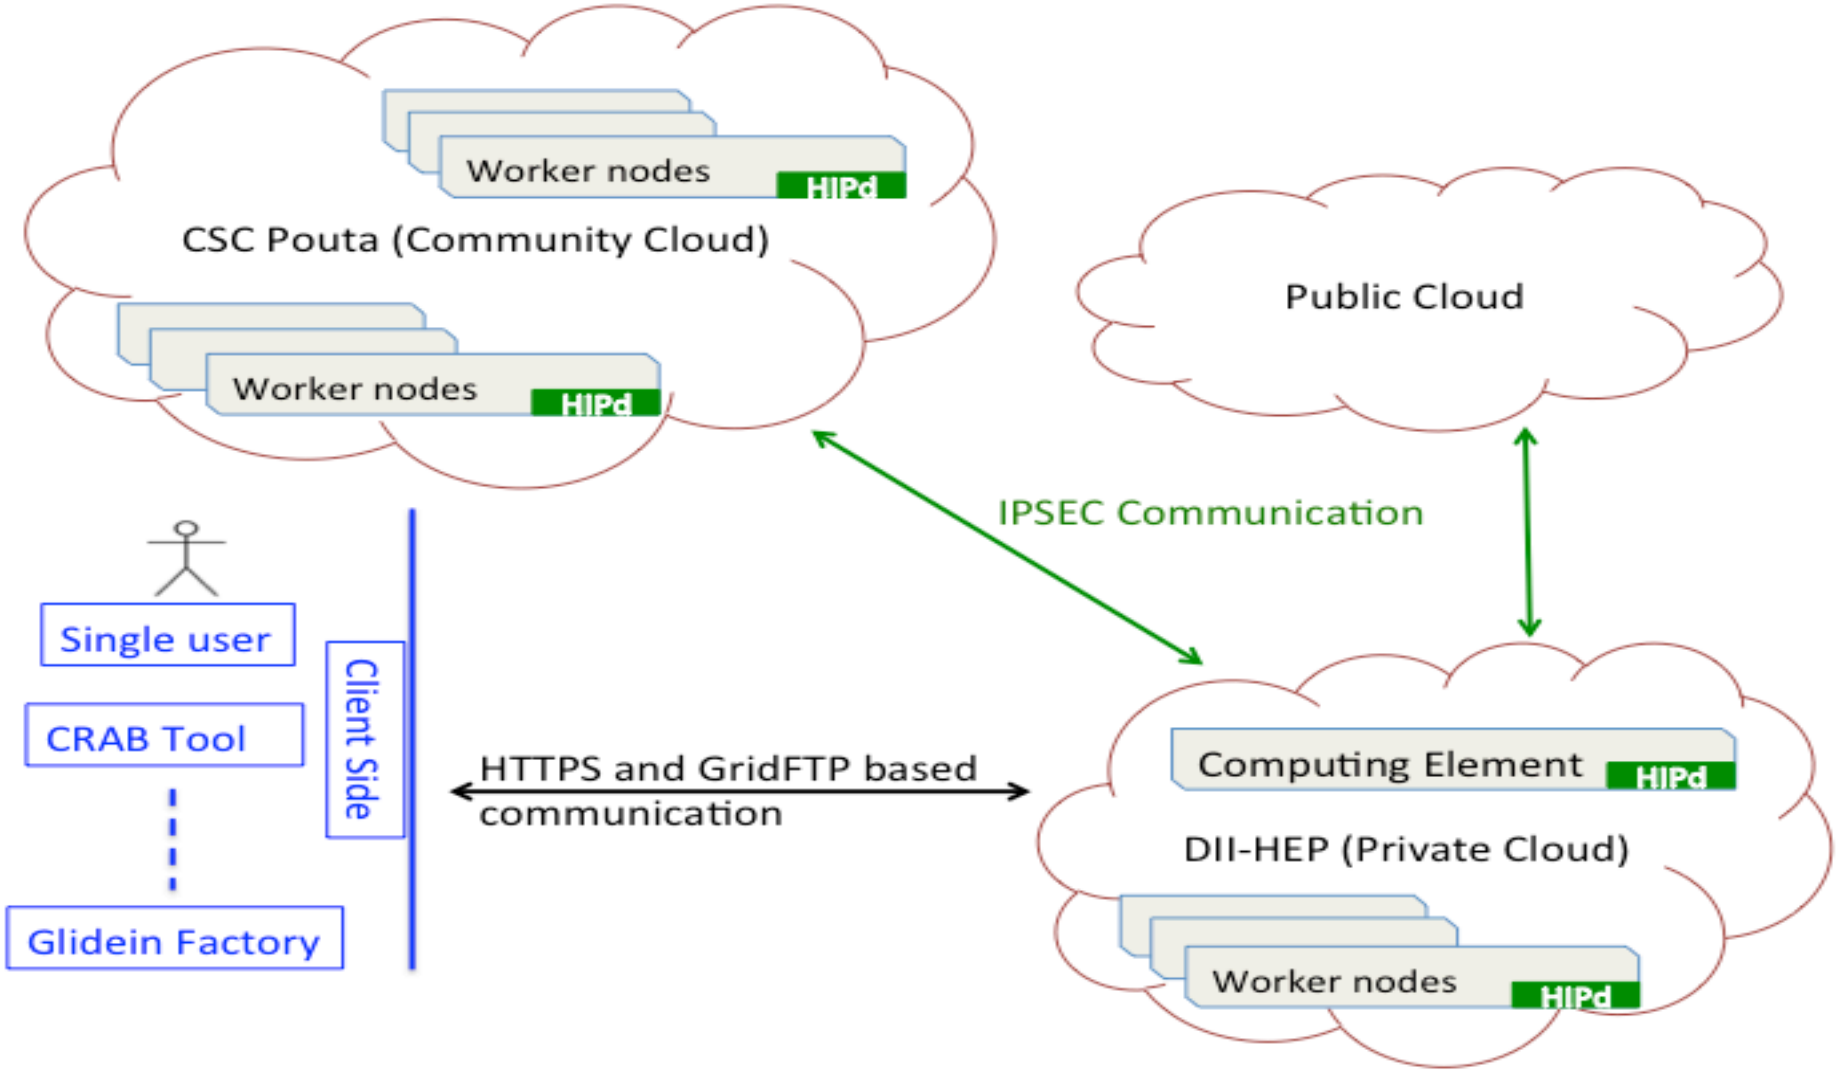

HIP based experiments

- Submitted 10,000 CMS jobs using CRAB tool
- Execution time is round 170 minutes
- Jobs are CPU bound but required data to be accessed via from external storage via *xrootd* Protocol
- Job success rate is 99%

Single core CPU utilization for CMS jobs

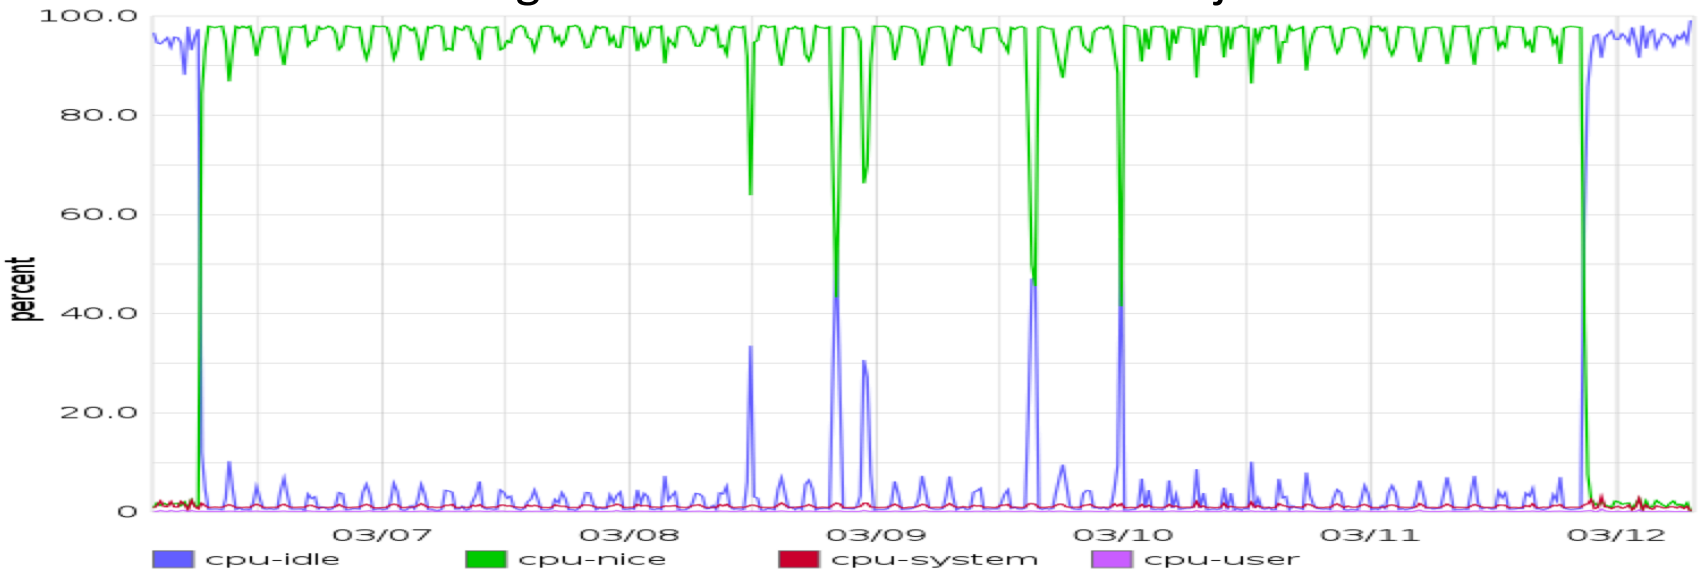

Computing Element (CE) CPU usage, without security and with HIP protocol

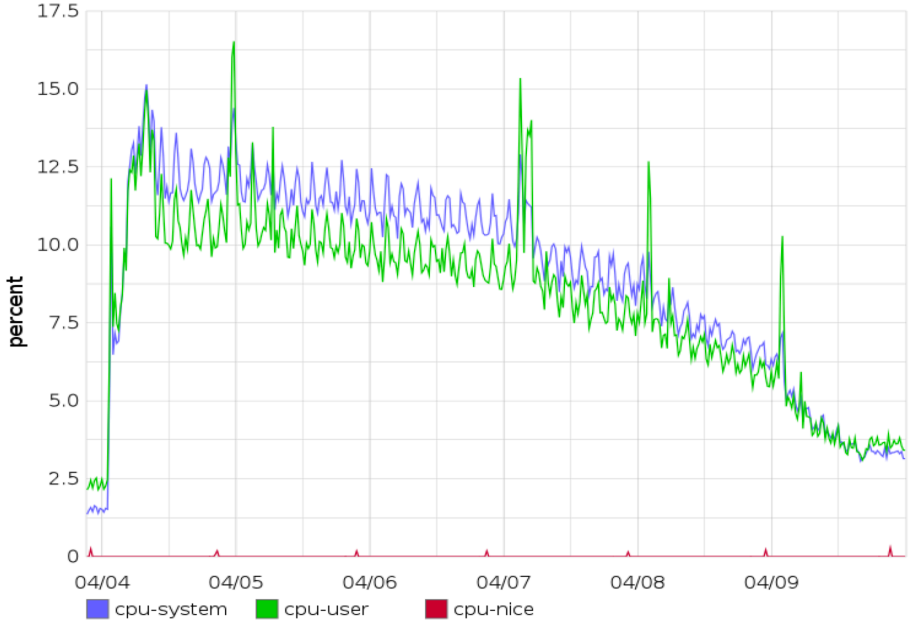

Computing Element (CE) Network usage, without security and with HIP protocol

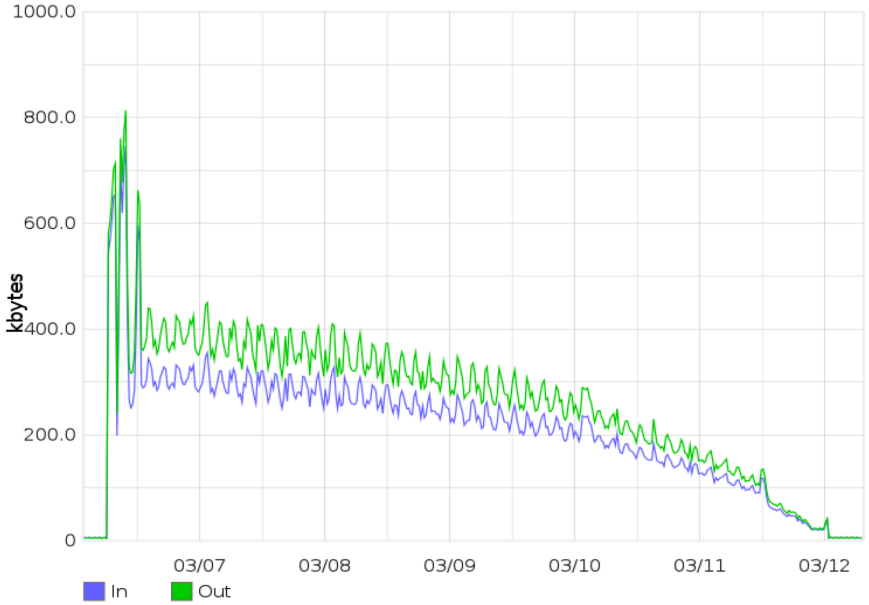

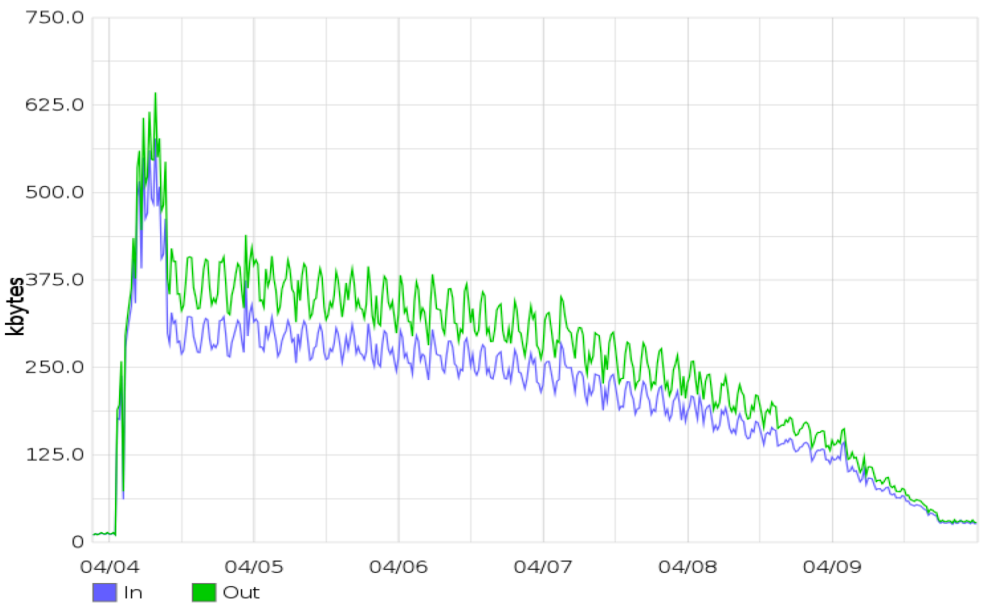

Summed CPU usage for Worker nodes, without security and with HIP protocol

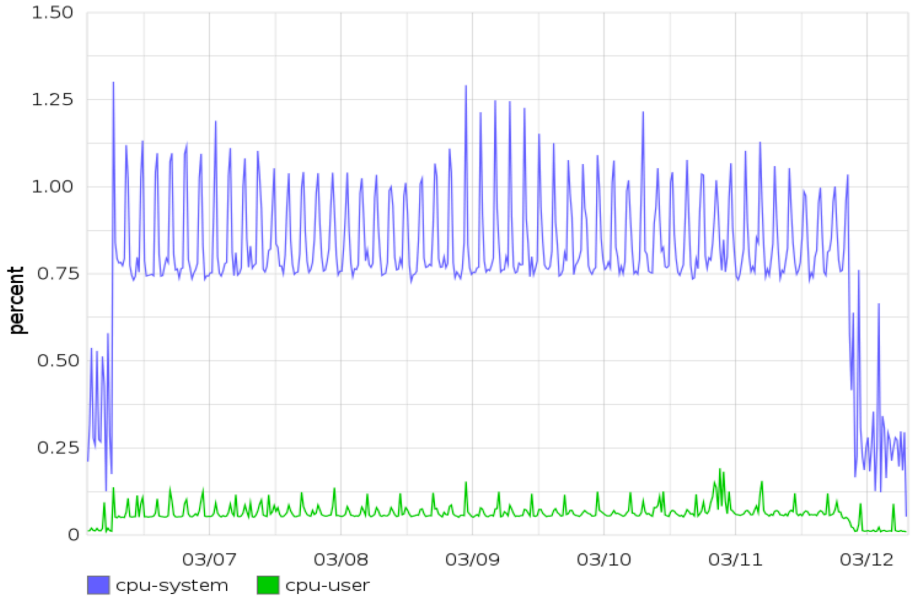

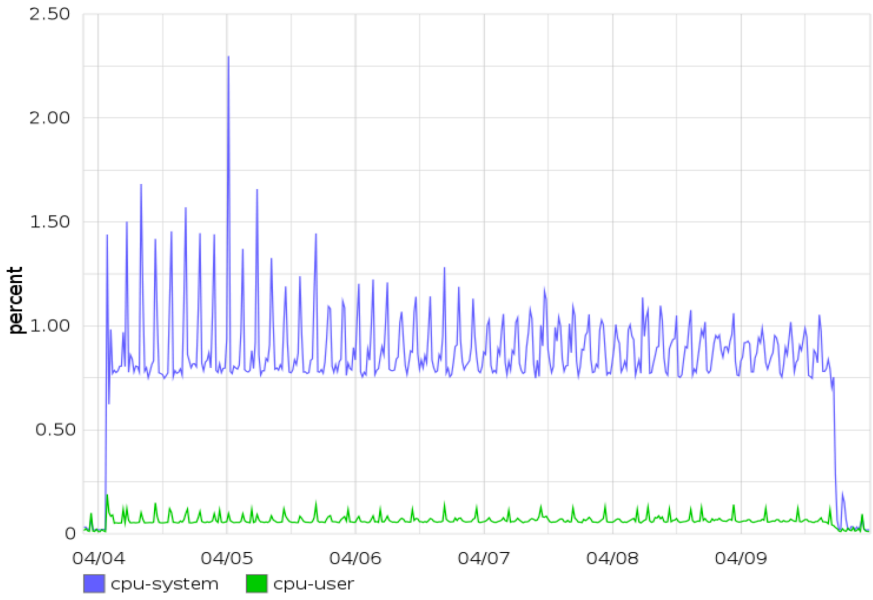

Summed Network usage for Worker nodes, without security and with HIP protocol

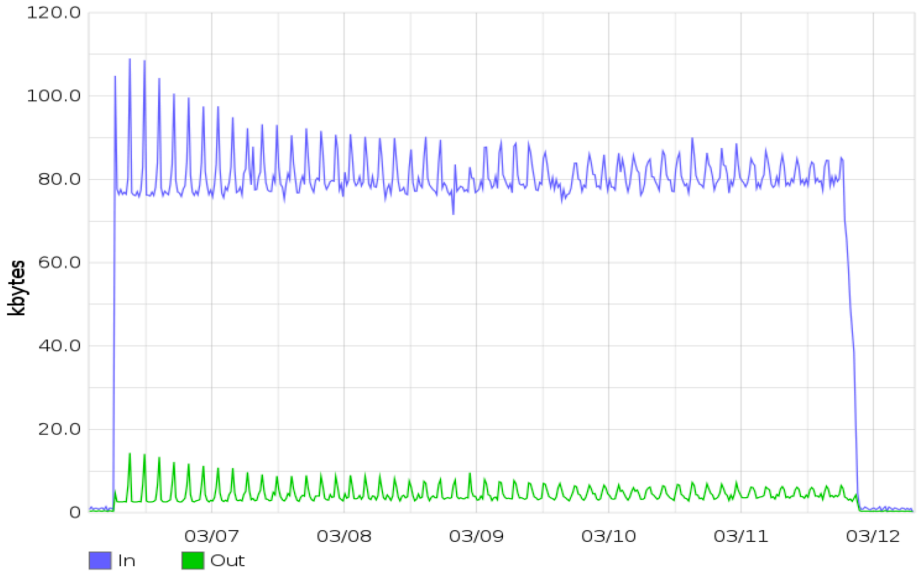

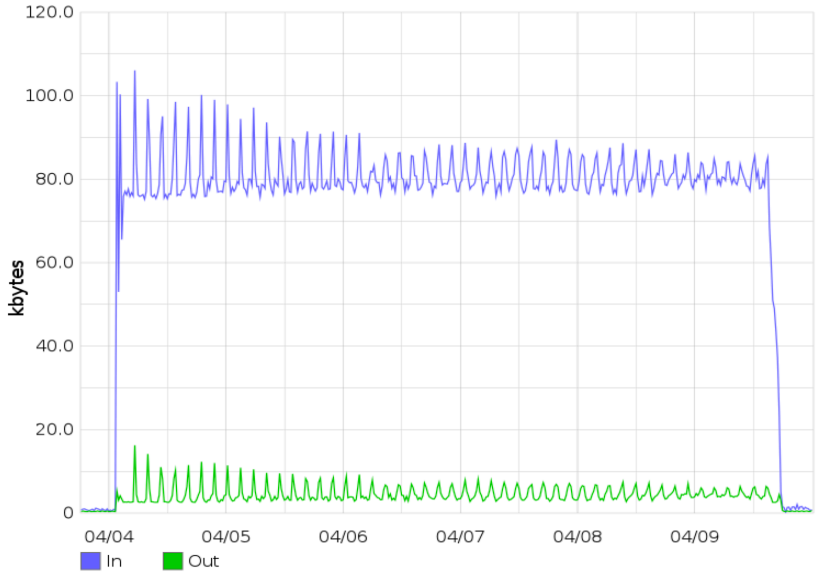

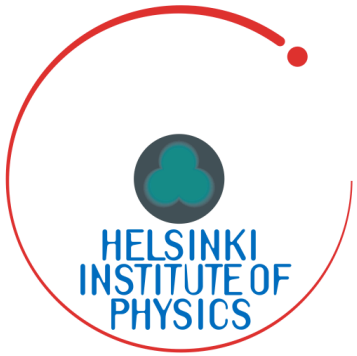

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

Article submitted in IEEE:

L. Osmani, S. Toor, M. Komu, M. J. Kortelainen, T. Lindén, J. White, R. Khan, P. Eerola, S. Tarkoma. **Secure Cloud Connectivity for Scientific Applications**. Submitted in IEEE Transactions on Cloud Computing, Special Issue on Scientific Cloud Computing.

Summary

- A scalable infrastructure for CMS jobs has been created which is cloud based and grid enabled
- The cloud setup itself does not rely on grid tools
- The same setup could be used for other applications as well
- The cloud provides flexibility in managing the infrastructure
- Initial results show that the HIP protocol can be very helpful in
  - secure multi-tenant environment
  - secure federated clouds
  - resource mobility

Future directions

- Work on making the EES plugin for OpenStack production quality is continuing
- The Jade cluster is planned to be replaced by a cloud setup on Pouta
- This cloud setup could be used by other scientific applications, like in the Finnish Grid and cloud Infrastructure consortium
- Planning to use Docker containers
- Also working on designing models for power measurements

Possibility of future collaboration

We will be happy to define new collaborations if there will be some common interests.

Prof. Sasu Tarkoma,
sasu.tarkoma@cs.helsinki.fi
http://www.cs.helsinki.fi/u/starkoma/

Dr. Salman Toor
salman.toor@helsinki.fi
http://www.it.uu.se/katalog/salto690

Dr. John White
john.white@cern.ch
http://hiptek.web.cern.ch/hiptek/Personnel/persons/person_0008.html

DII-HEP Cloud

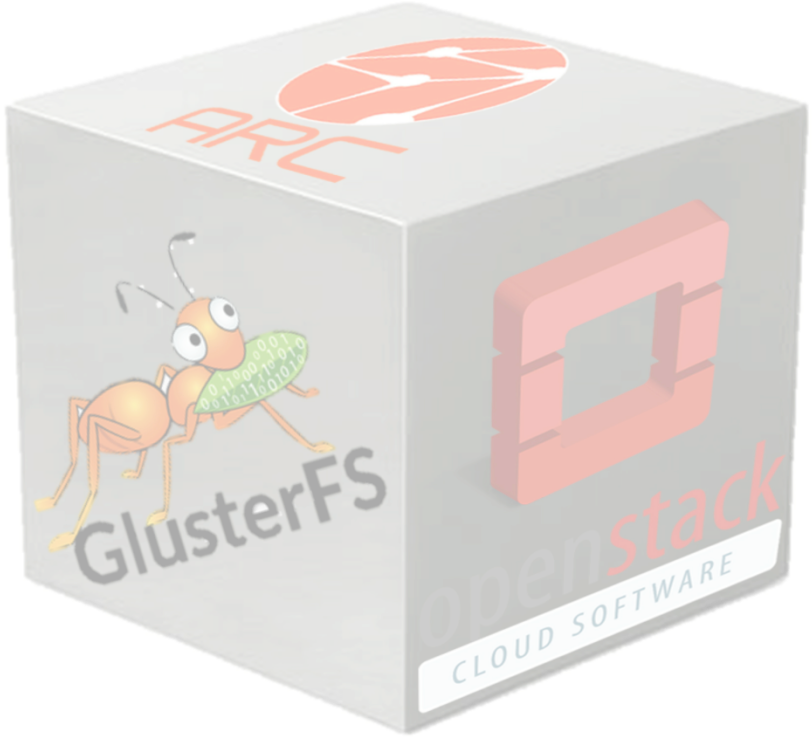

Thanks!