Particle Physics with Cloud Computing

Tomasz Malkiewicz

CSC

Outline

- Computational particle physics
- Where we are
  - – HPC
  - – Grid
- Where we are heading
  - – Cloud
- Cloud computing in Finland
  - – cPouta
    - Particle physics related projects
  - – Ongoing/coming project

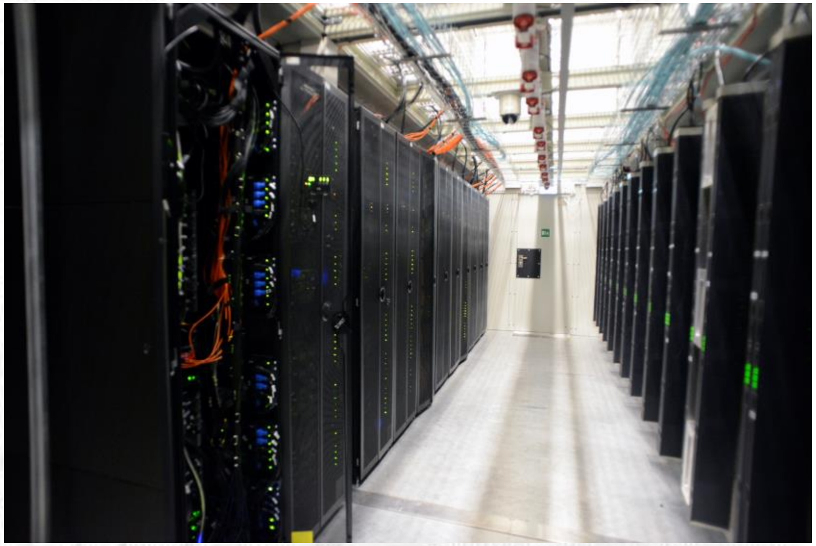

Computational particle physics

- *Methods* and *computing tools* developed in and used by particle physics research
- Main branches
  – Lattice field theory (numerical computations)
  – Calculations of particle interaction or decay (computer algebra)
  – Event generators (stochastic methods)

Supercomputing: serial and parallel processing

- Serial computing
  - single processing unit (core) is used for solving a problem
  - single task performed at once
- Parallel computing
  - multiple cores are used for solving a problem
  - problem is split into smaller subtasks
  - multiple subtasks are performed simultaneously

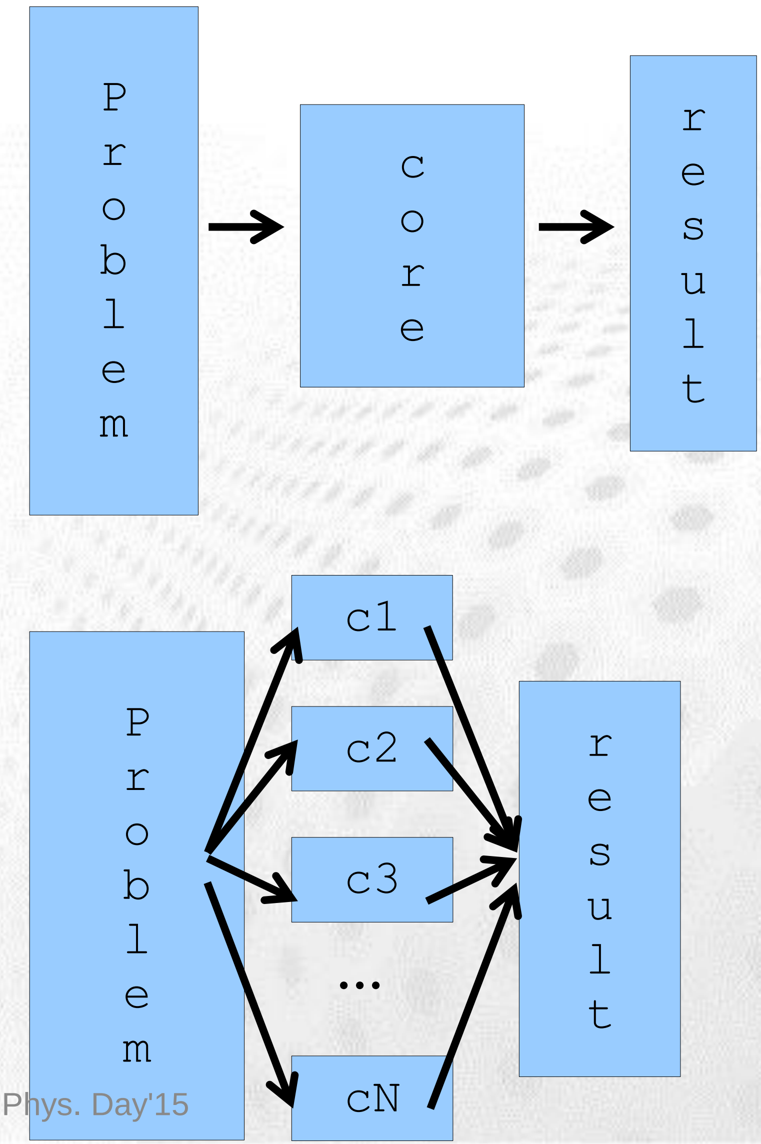

Data parallelism

- Data is distributed to processor cores
- Each core performs (nearly) identical tasks with different data
- Example: summing the elements of a 2D array

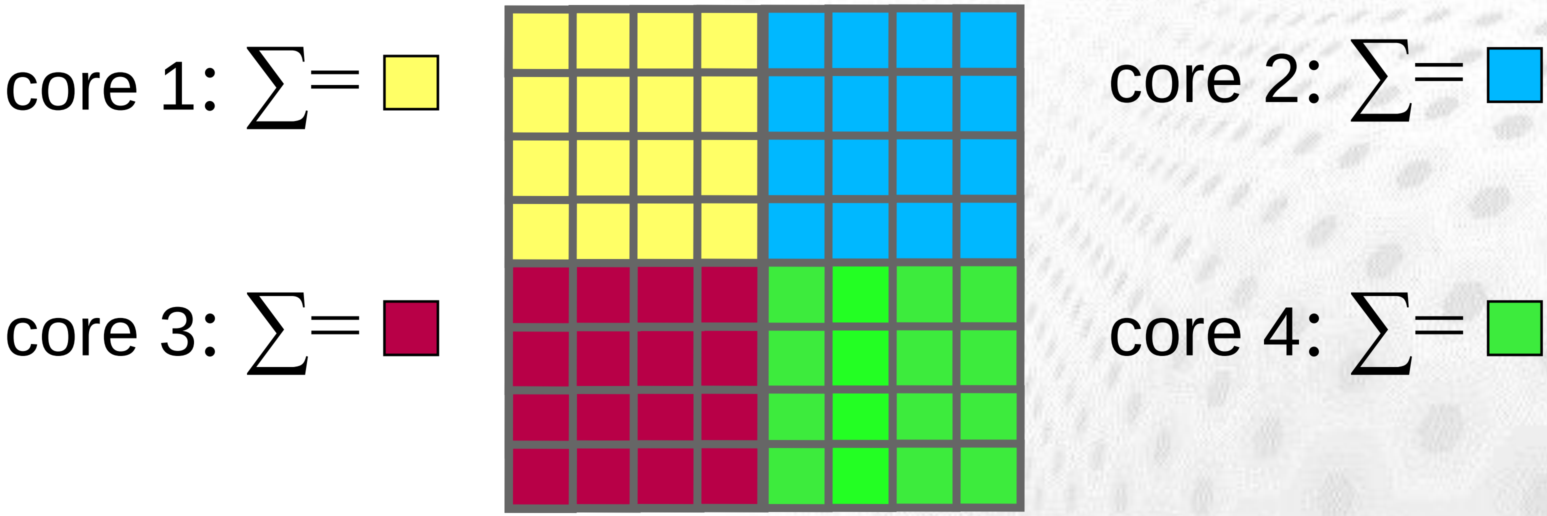

- Each core sums it's part of the array
- The individual sums have to be combined in the end

Task parallelism

- Different cores perform different tasks with the same or different data
- Example: signal processing, four filters as separate tasks

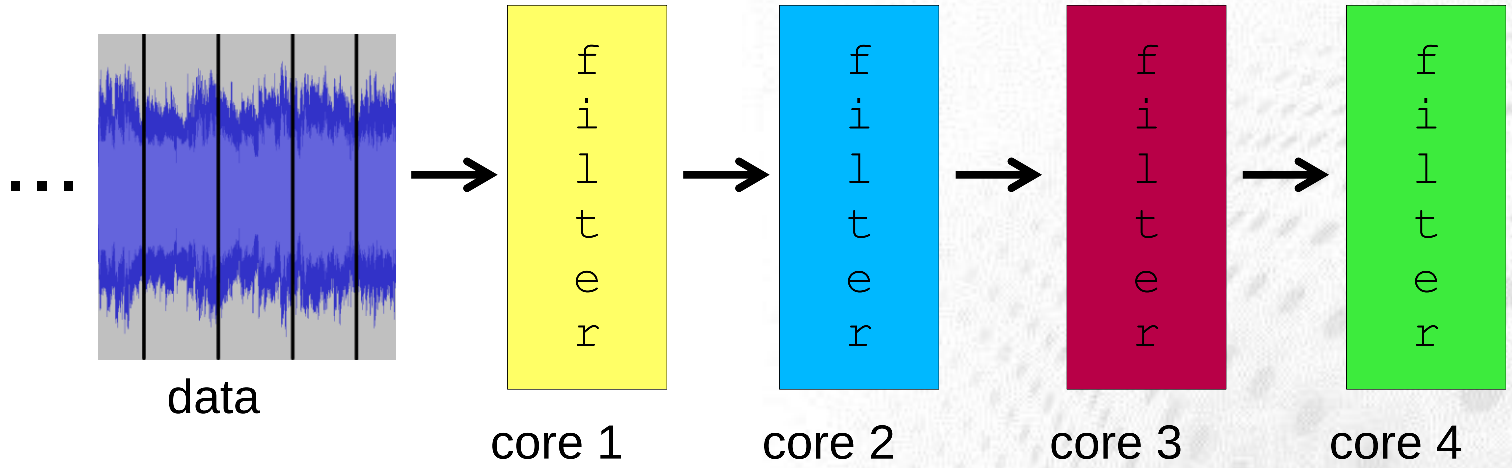

- Data is processed as segments
- Core 2 obtains a segment after core 1 has processed it; core 1 starts to process a new segment
- When the first segment gets to core 4, all cores are busy

Where we are

Finnish computing resources concentrated in Kajaani

Cray Dragonfly Topology

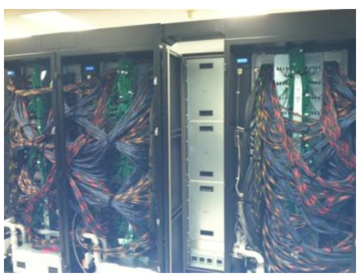

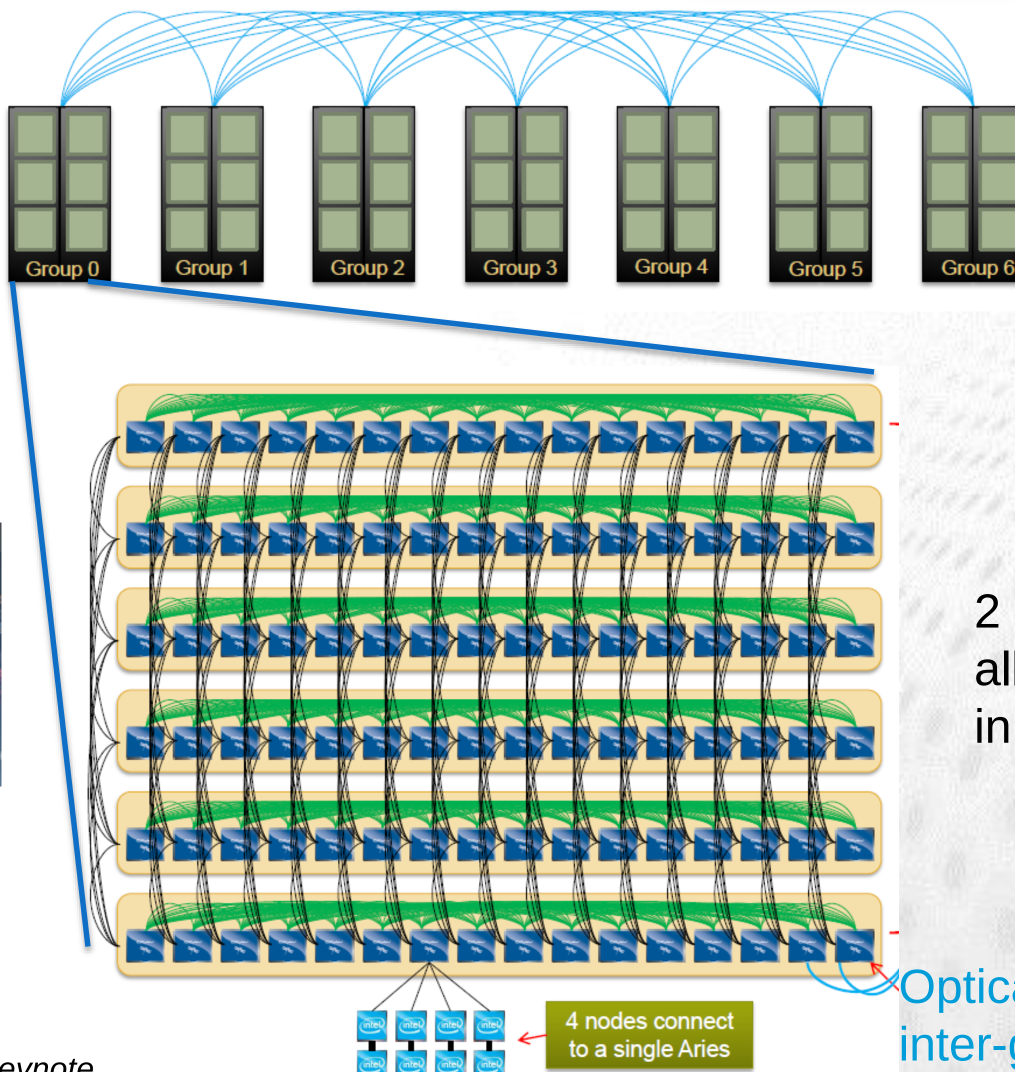

All-to-all network between groups

2 dimensional all-to-all network in a group

Optical uplinks to inter-group net

Source:
***Robert Alverson**, Cray*
Hot Interconnects 2012 keynote

Sisu rear view

Performance of numerical libraries

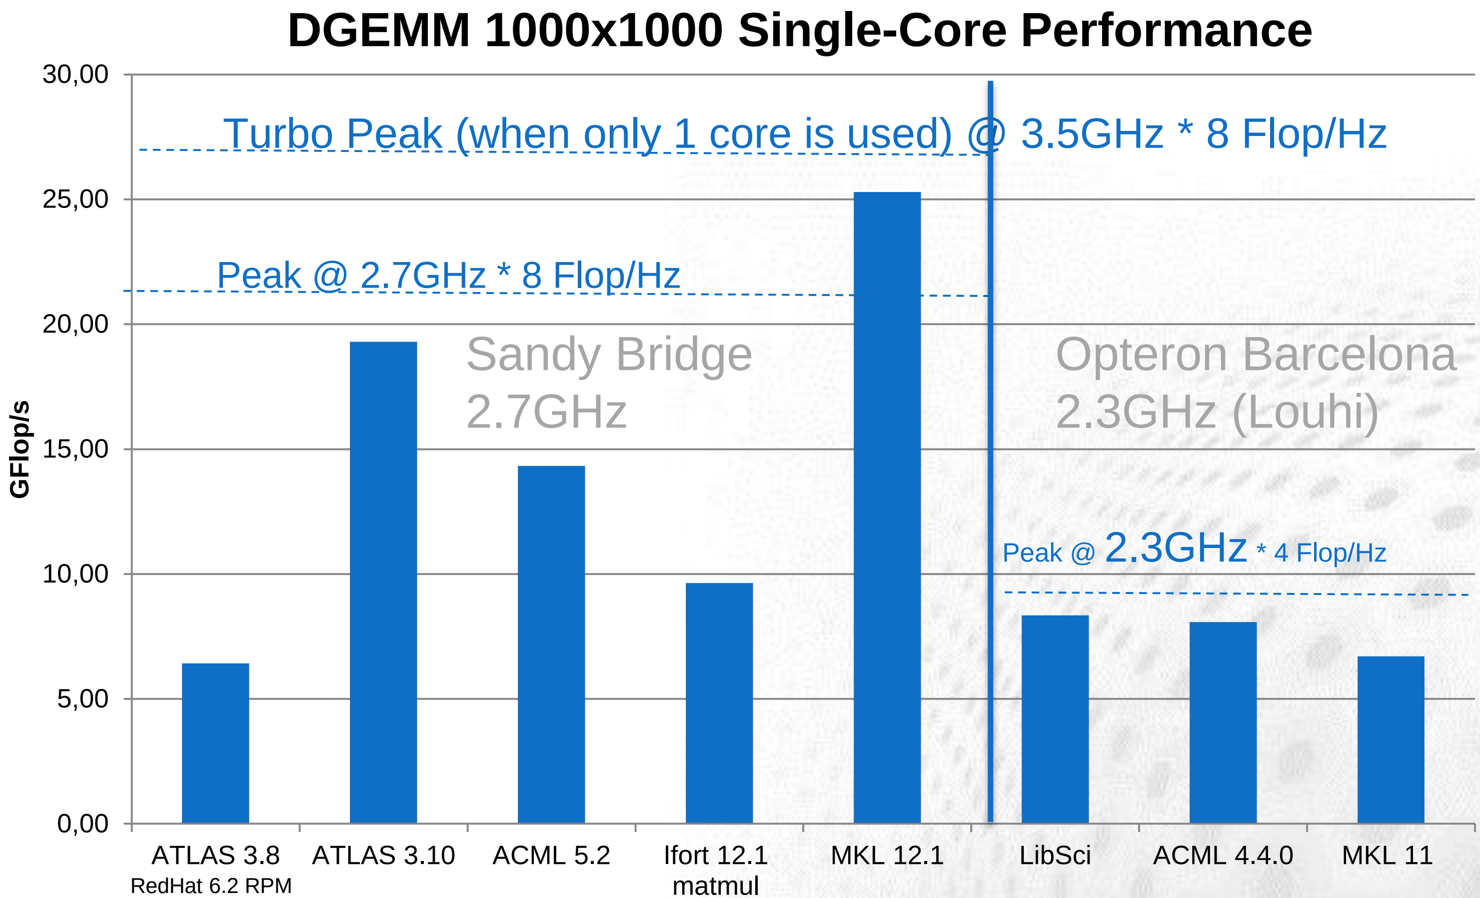

MKL the best choice on Sandy Bridge

(On Cray, LibSci a good alternative)

Emerging technologies – CUDA6

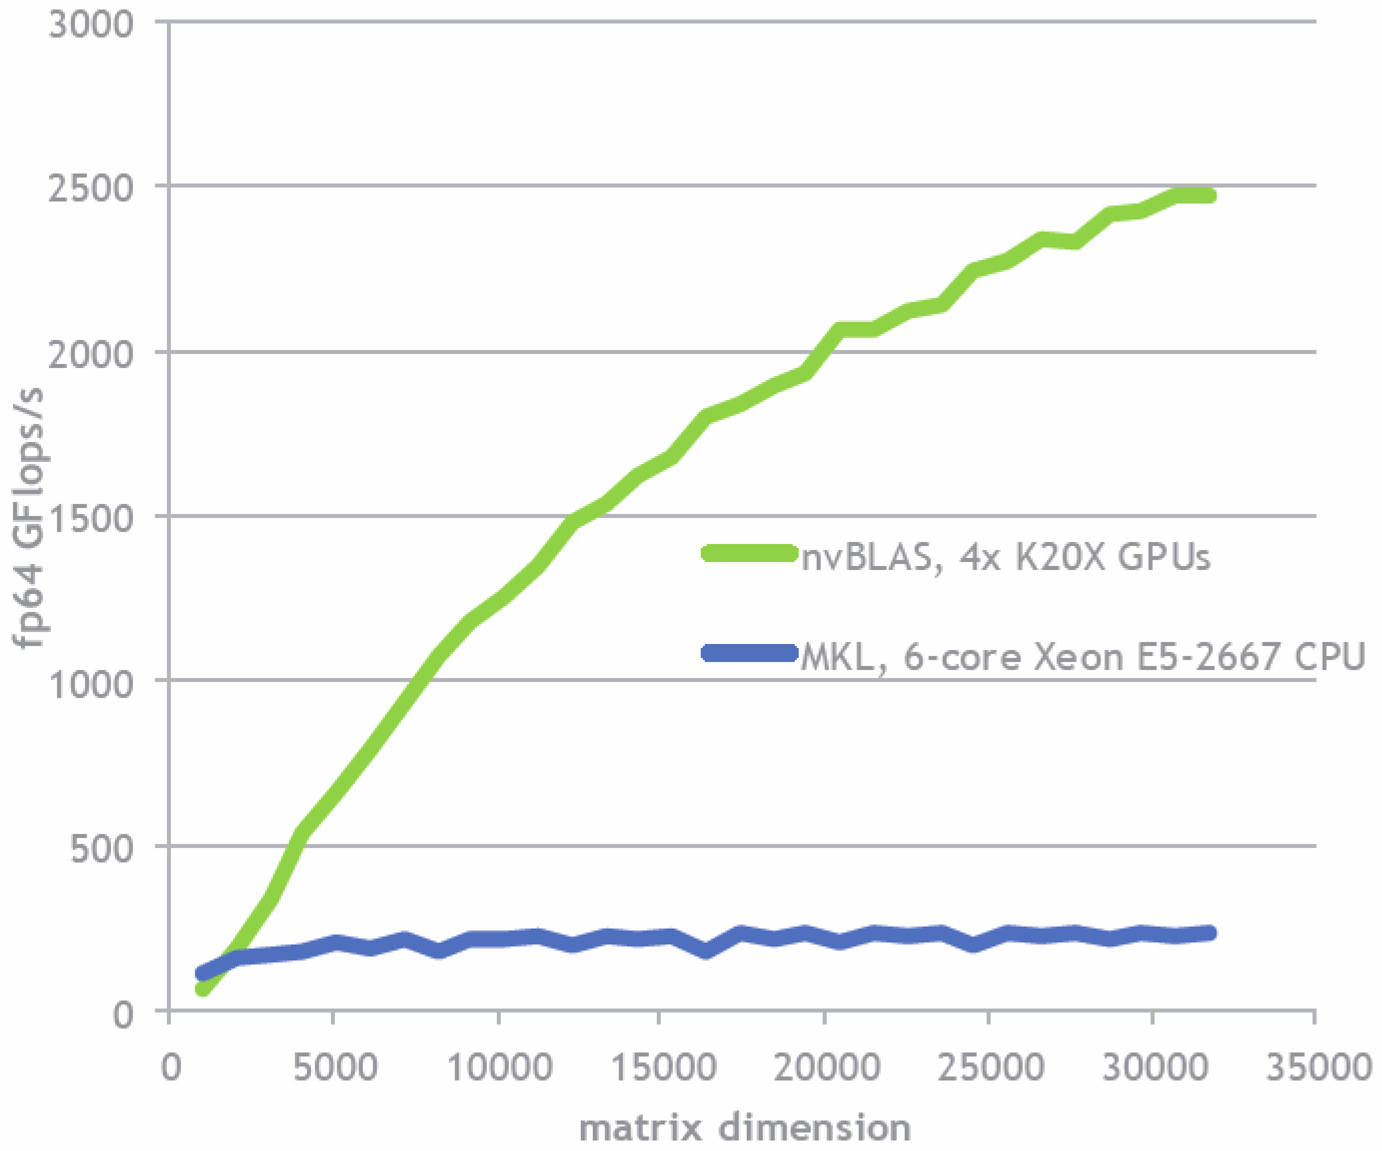

Source:
Peter Messmer,
CUDA lecture @CSC, 2014

Emerging technologies – Xeon Phi

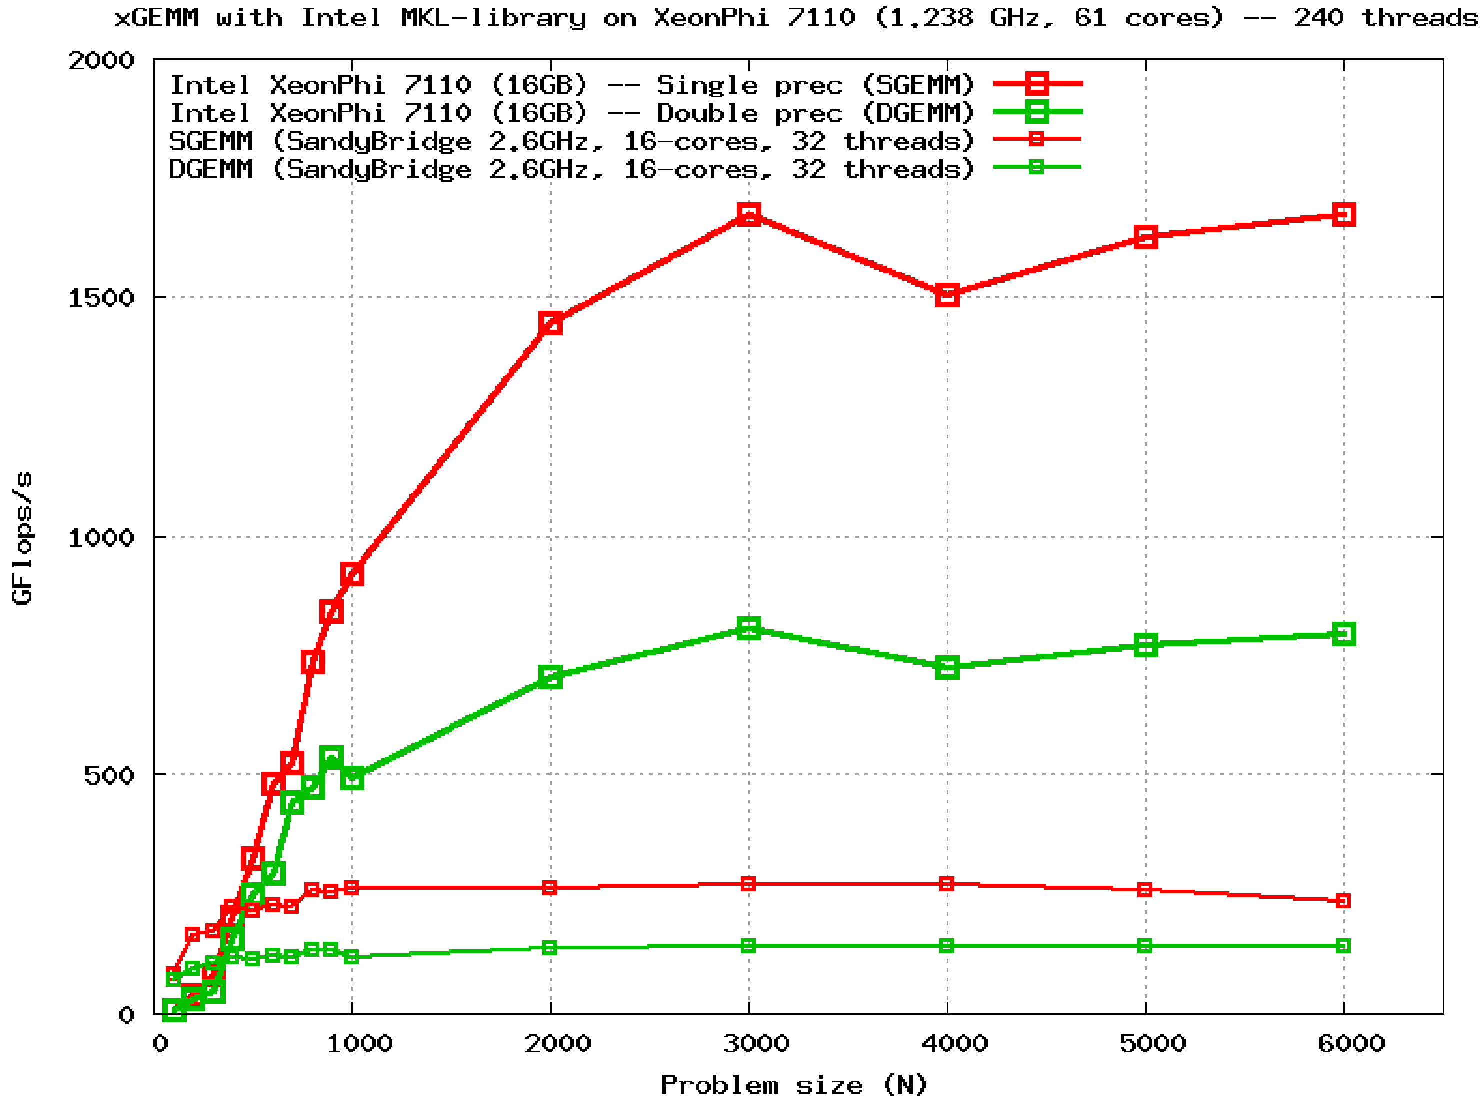

Use case 1: Solving quantum field theories

Lattice simulations:

Model the quantum field on a discrete lattice and use computer simulations:

- Strong coupling (QCD, technicolor, etc)
- Phase transitions
- High temperature and/or density
- Non-linear physics (Real-time evolution:
- Heavy ions/Cosmology

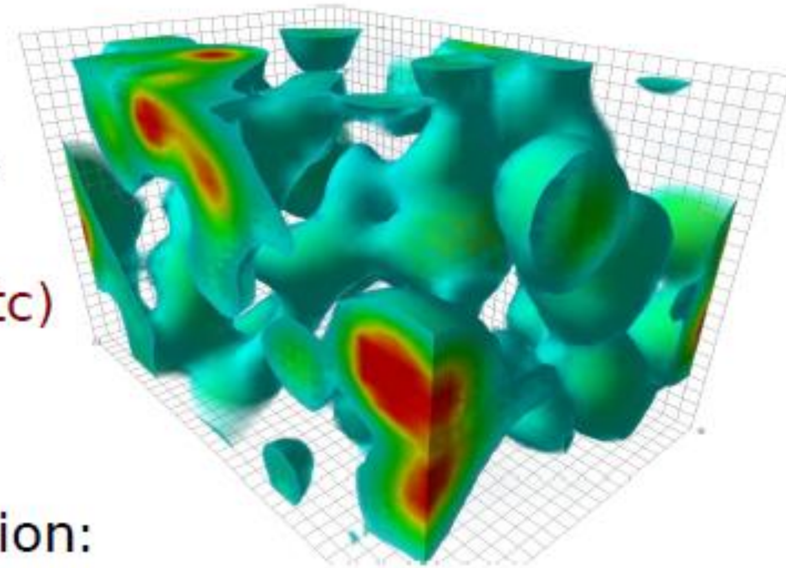

Finland: long traditions, 1st high-T QCD simulation in 1981 (Kajantie, Montonen, Pietarinen)

*Source: **Kari Rummukainen**, CSC Autumn School 2013*

Use case 1 cont.: Why simulations are expensive

CSC

- 4-dimensions (in some cases 3)
- quantum
- fermions (quarks)

Fermions can be integrated out, but that makes the action *non-local*: *fermion determinant*

$$Z = \int dU\, d\psi\, e^{-S_G(U) - \bar{\psi} M(U) \psi} = \int dU\, det[M(U)] e^{-S_G(U)}$$

$$= \int dU\, d\chi\, e^{-S_G(U) - \bar{\chi} M(U)^{-1} \chi}$$

last step: use commuting *pseudofermion* fields: det → inverse

Pseudofermions: normal complex numbers

Expensive part: $M^{-1}(U)\,\chi$ - *conjugate gradient*

matrix size: $[\text{volume} \times 4 \times 3]^2$

If volume 64^4, matrix is $[201\,326\,592]^2$

Source: ***Kari Rummukainen****,*
CSC Autumn School 2013

Use case 1 cont.: Parallel paradigms

- Pure MPI: 1 process/local volume
  - relatively simple
  - local volume can be small if communication is efficient & async
  - e.g. 3^4 local volume, 48^4 total volume, 65536 cores (BlueGene/Q)
- OpenMP (multithreaded) without MPI
  - Easiest, close to non-parallel code
  - Neeeds independent loop iterations – “vectorized code”
  - Efficient only in shared memory – only 10's of cores
- Mixed mode: MPI + within local volume OpenMP / thread pool
  - Potentially highest performance, because can use more cores/volume
  - Perf/core is usually not better than pure MPI
  - In sisu: 16 cores/node, 1-16 threads/MPI process
  - Blue Gene Q: 32-64 threads/node

Source: ***Kari Rummukainen****,*
CSC Autumn School 2013

Use case 2: LENA simulations

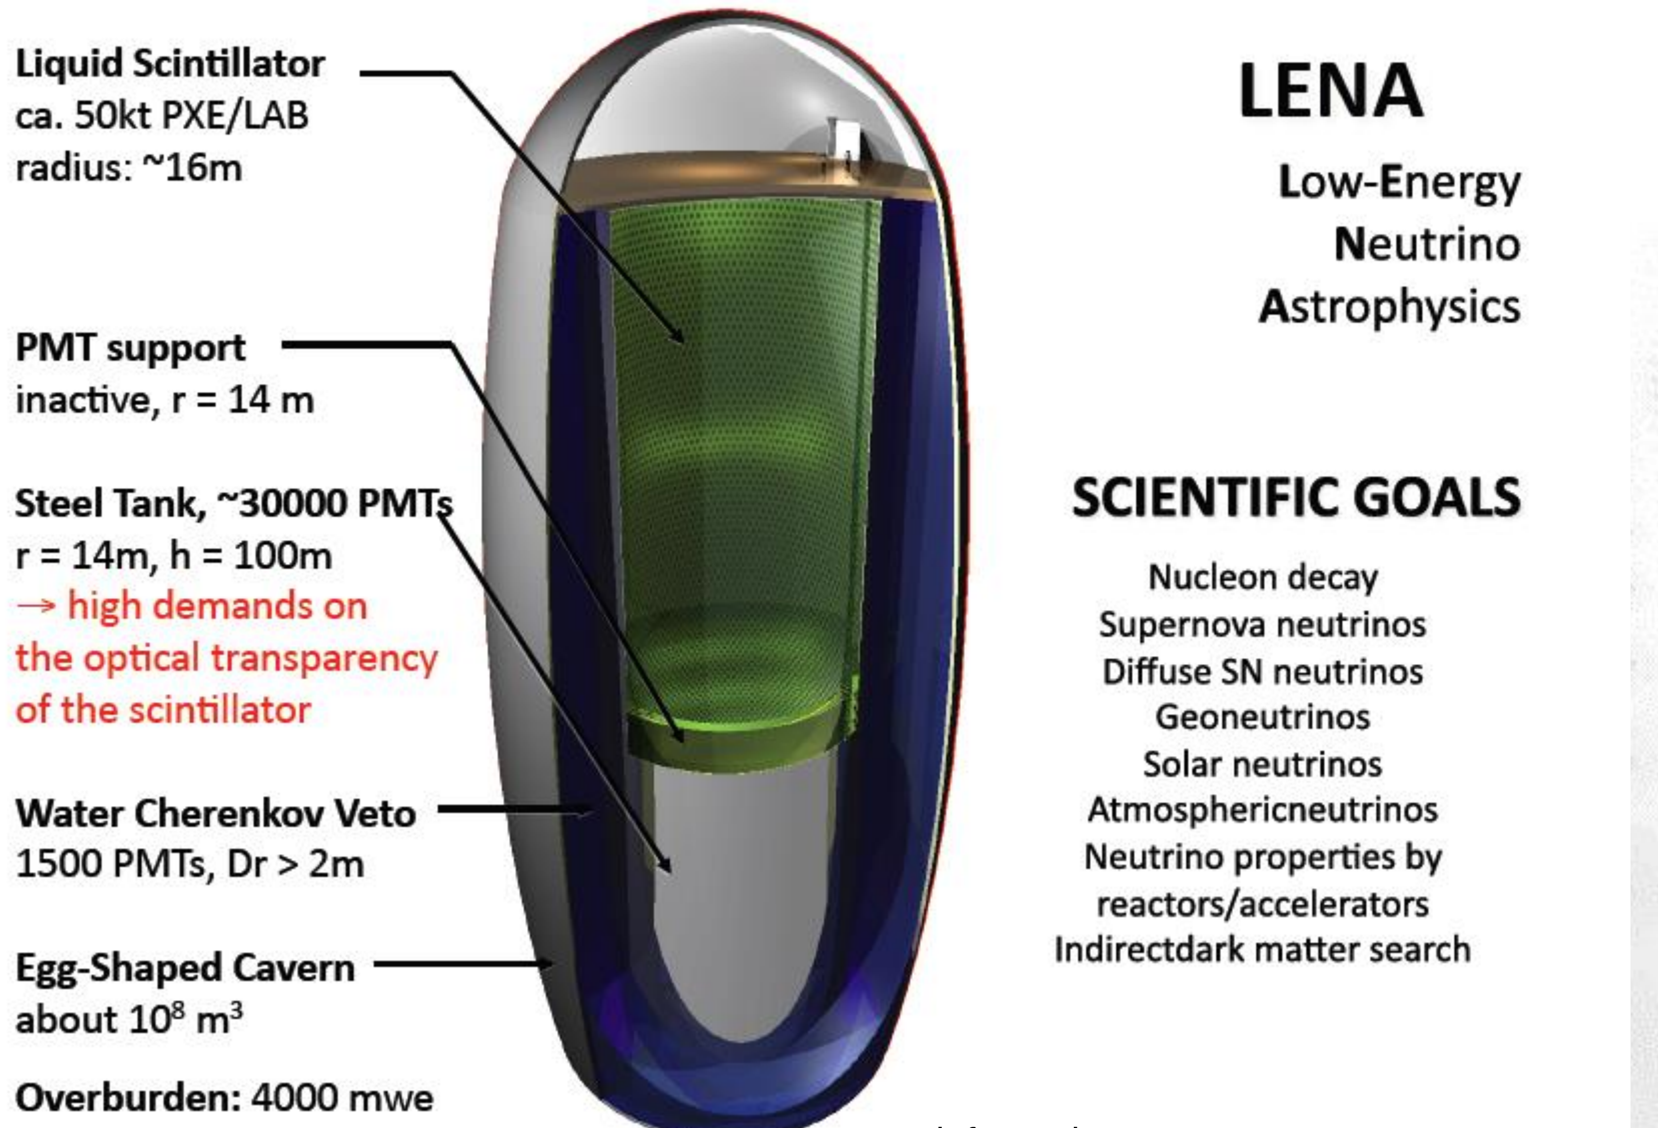

Source: ***Kai Loo****, HPC Nuclear and Particle Physics seminar at CSC 2012*

More information:

Poster no. 66 on Poster session II by K. Loo

Where we are heading

Cloud Service Models

SaaS

Software: End user applications,
Not only web applications.
Integrated also e.g. file system view, job submission.

PaaS

Platform: Virtual Machine images with *preinstalled software* e.g. OS, DB, Hadoop, cluster tools, code development tools, science discipline tailored applications.

IaaS

Infrastructure: Cloud user interface to launch Virtual Machines where the user can choose operating system and administrate it.

Compute resources: CPU, RAM
Storage: local disks or via network (NFS, http,…)
Network: internal, external

Cloud Computing in Finland

cCloud running on Taito supercluster

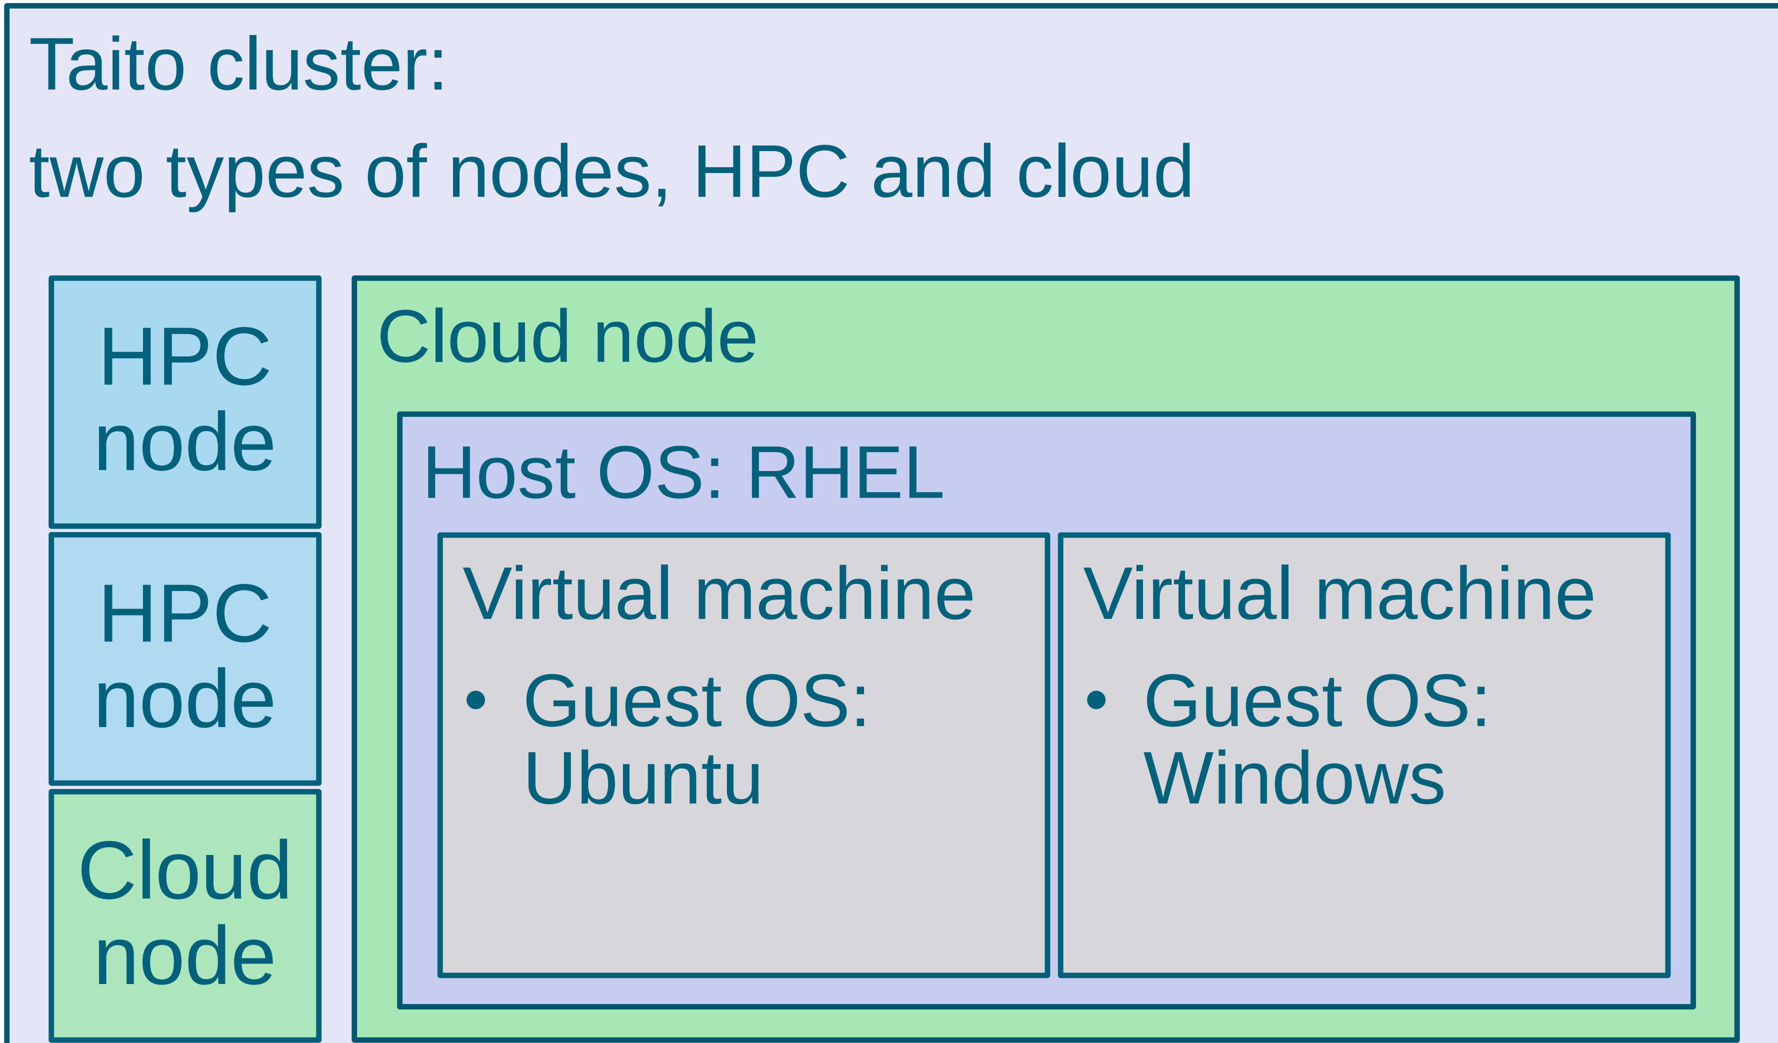

cPouta

- Infrastructure as a Service (IaaS)
  – Deploy your own virtual machines/storage/network
- OpenStack cloud middleware
  – To manage the virtual resources
- Allocated from Taito supercluster
  – Powerful CPUs, lots of memory, fast disk
- Simple to create and modify VMs
  – WWW, CLI & REST API interfaces

https://research.csc.fi/pouta-iaas-cloud

Specifications of the resources

	Cores	Memory	Disk (root)	Disk (ephemeral)	Disk (total)	Memory/core	Billing Units/h
tiny	1	1 GB	10 GB	110 GB	120 GB	1	2
mini	1	3,5 GB	10 GB	110 GB	120 GB	3	2
small	4	15 GB	10 GB	220 GB	230 GB	4	8
medium	8	30 GB	10 GB	440 GB	450 GB	4	16
large	12	45 GB	10 GB	660 GB	670 GB	4	24
bigroot	16	60 GB	80 GB	500 GB	580 GB	4	32
fullnode	16	60 GB	10 GB	900 GB	910 GB	4	32

- VM configuration: flavor. Project has total quota.
- Server-class Xeon E5 (Sandy Bridge) CPUs
  - 2*8 cores, 2.6GHz
- Connected with 40 Gbit Ethernet
- Block storage from DDN back-end
  - hard drive attached to a single VM

Division of work

IaaS cloud expert

- **Provides**
  - **Resources (compute, storage)**
  - **Interfaces** to access the system
- **Supports** usage of the cloud, but does not necessarily manage Virtual Machines (VM)
  - Does not know what is running on the VMs

VM admin

- **Can connect the existing** compute / storage **resources** through the **private network solution**
- **Manages** Virtual Machines
  - **`root` permission** for VMs
  - **Installs and maintains** Operating System and other software for VMs
  - Pays the software licenses

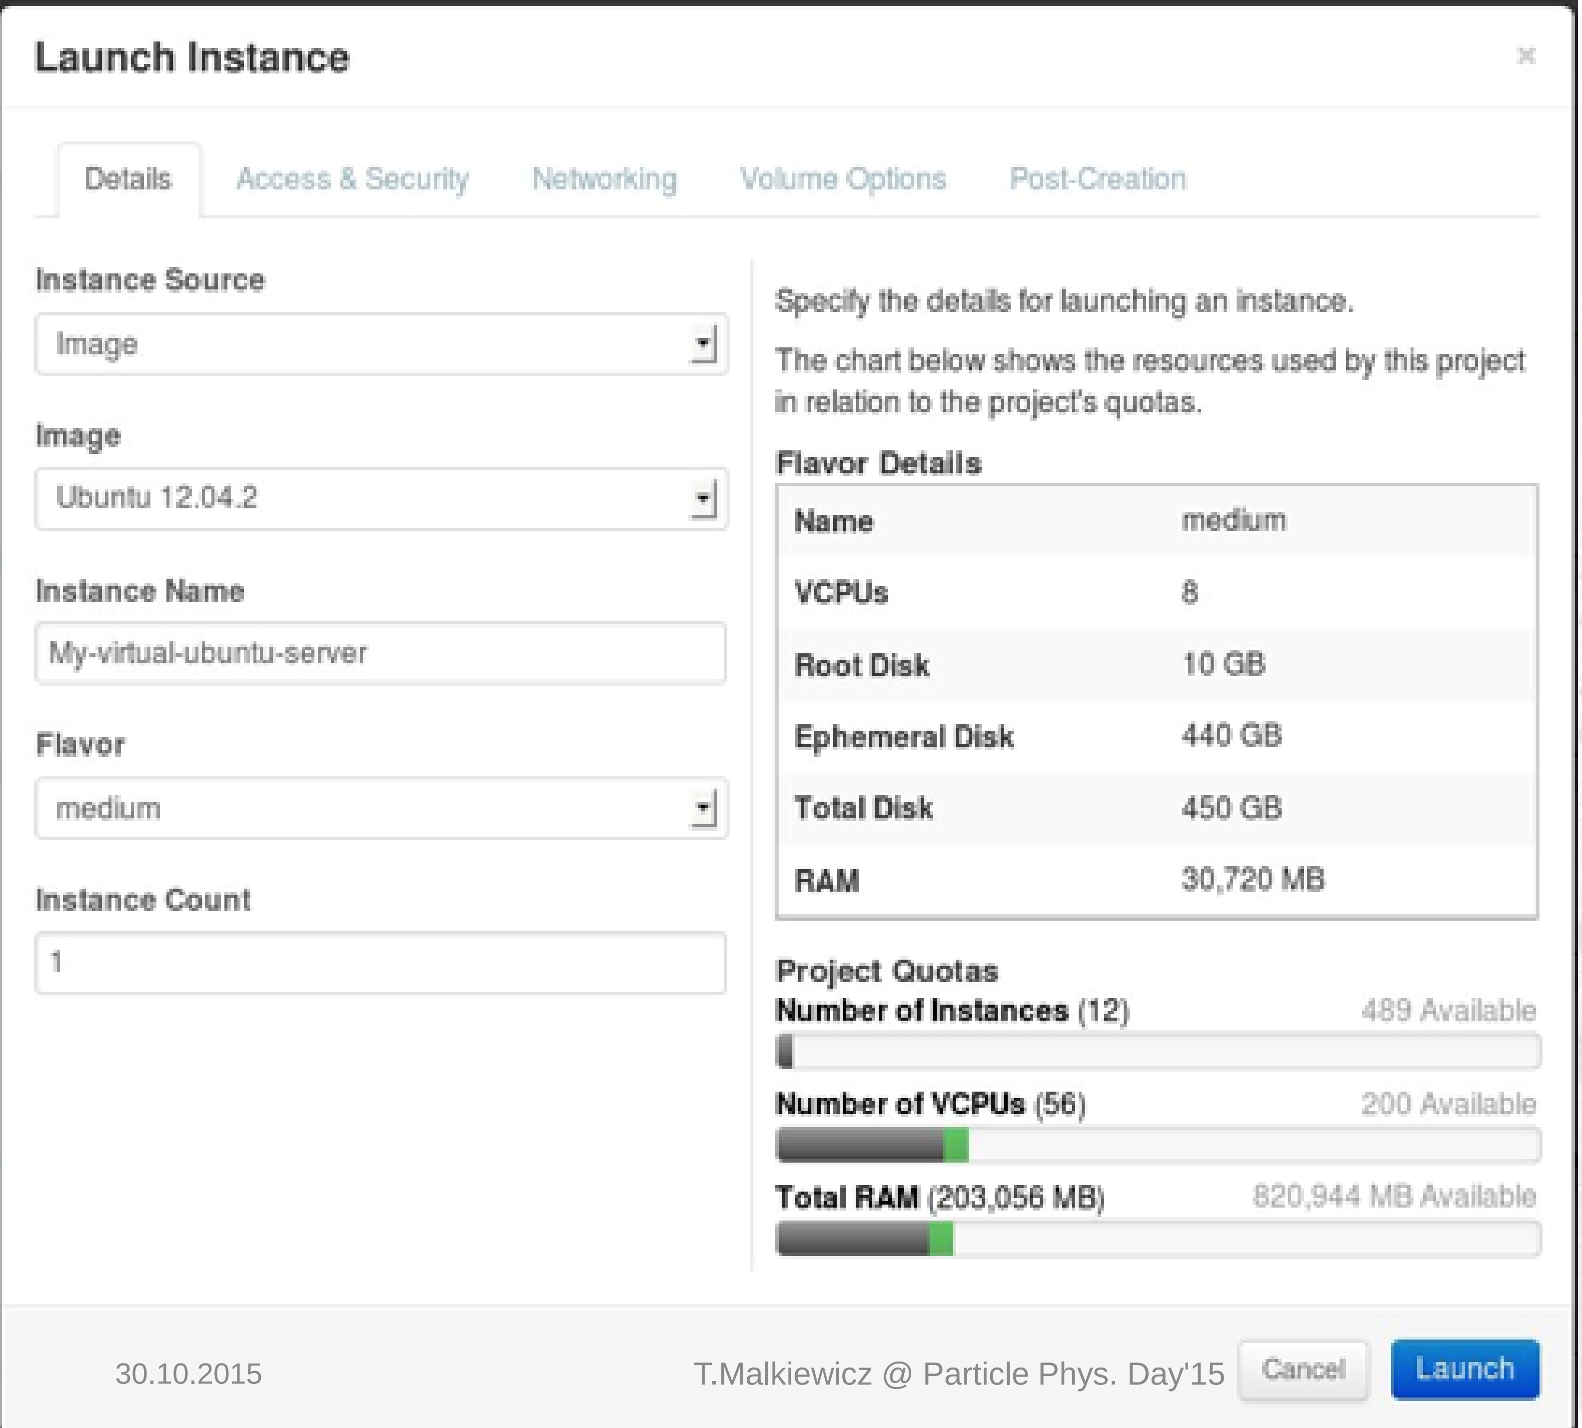
Launch Instance
Details
Access & Security
Networking
Volume Options
Post-Creation
Instance Source
Image
Image
Ubuntu 12.04.2
Instance Name
My-virtual-ubuntu-server
Flavor
medium
Instance Count
1
Specify the details for launching an instance.
The chart below shows the resources used by this project in relation to the project's quotas.
Flavor Details
Name medium
VCPUs 8
Root Disk 10 GB
Ephemeral Disk 440 GB
Total Disk 450 GB
RAM 30,720 MB
Project Quotas
Number of Instances (12) 489 Available
Number of VCPUs (56) 200 Available
Total RAM (203,056 MB) 820,944 MB Available
Cancel
Launch

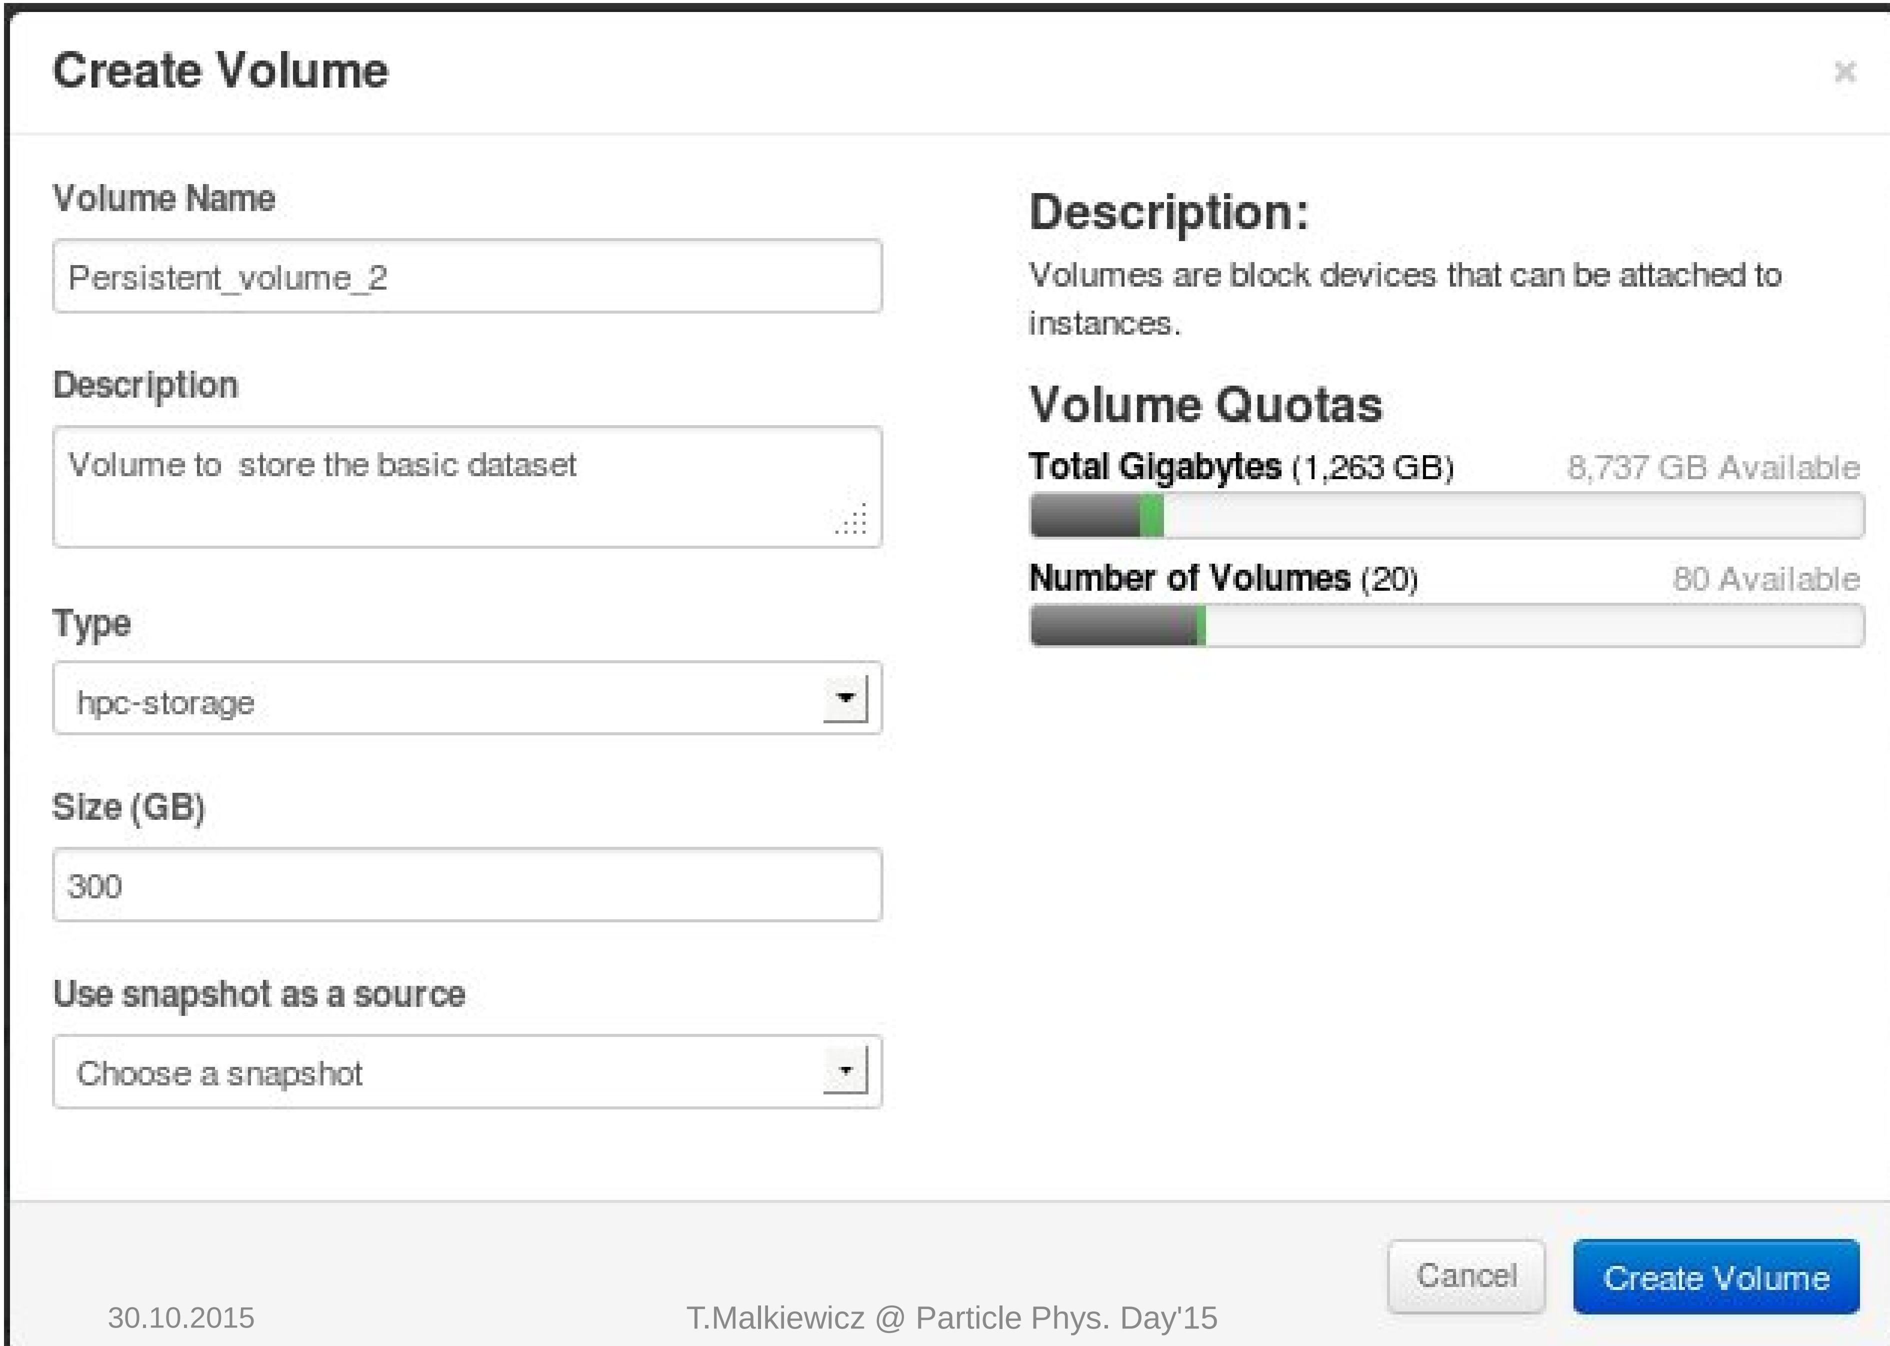
Create Volume
×
Volume Name
Persistent_volume_2
Description
Volume to store the basic dataset
Type
hpc-storage
Size (GB)
300
Use snapshot as a source
Choose a snapshot
Description:
Volumes are block devices that can be attached to instances.
Volume Quotas
Total Gigabytes (1,263 GB)
8,737 GB Available
Number of Volumes (20)
80 Available
Cancel
Create Volume

Ready to go

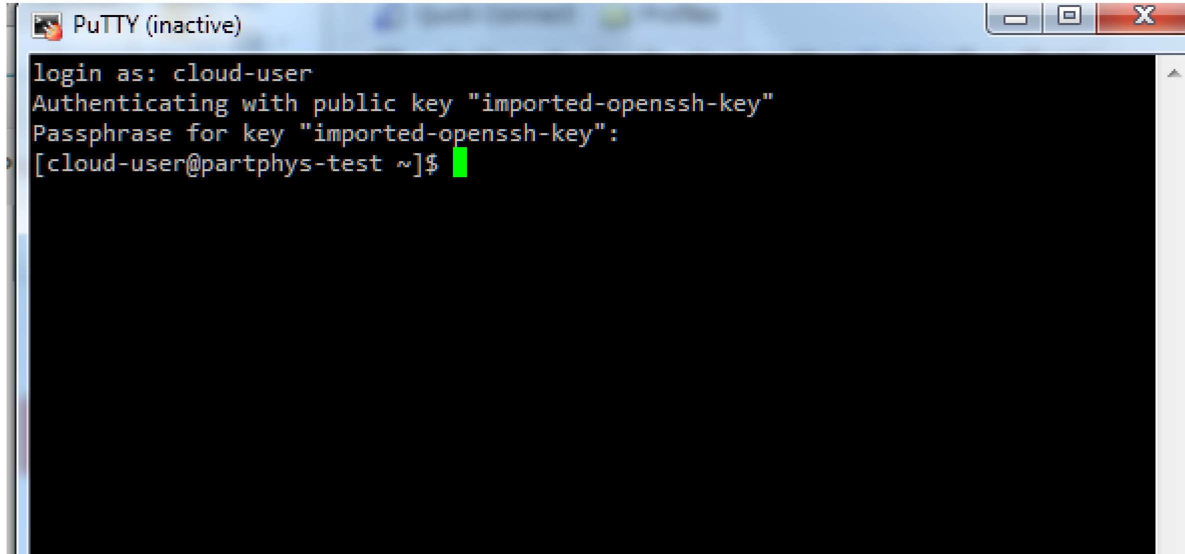
PuTTY (inactive)
login as: cloud-user
Authenticating with public key "imported-openssh-key"
Passphrase for key "imported-openssh-key":
[cloud-user@partphys-test ~]$

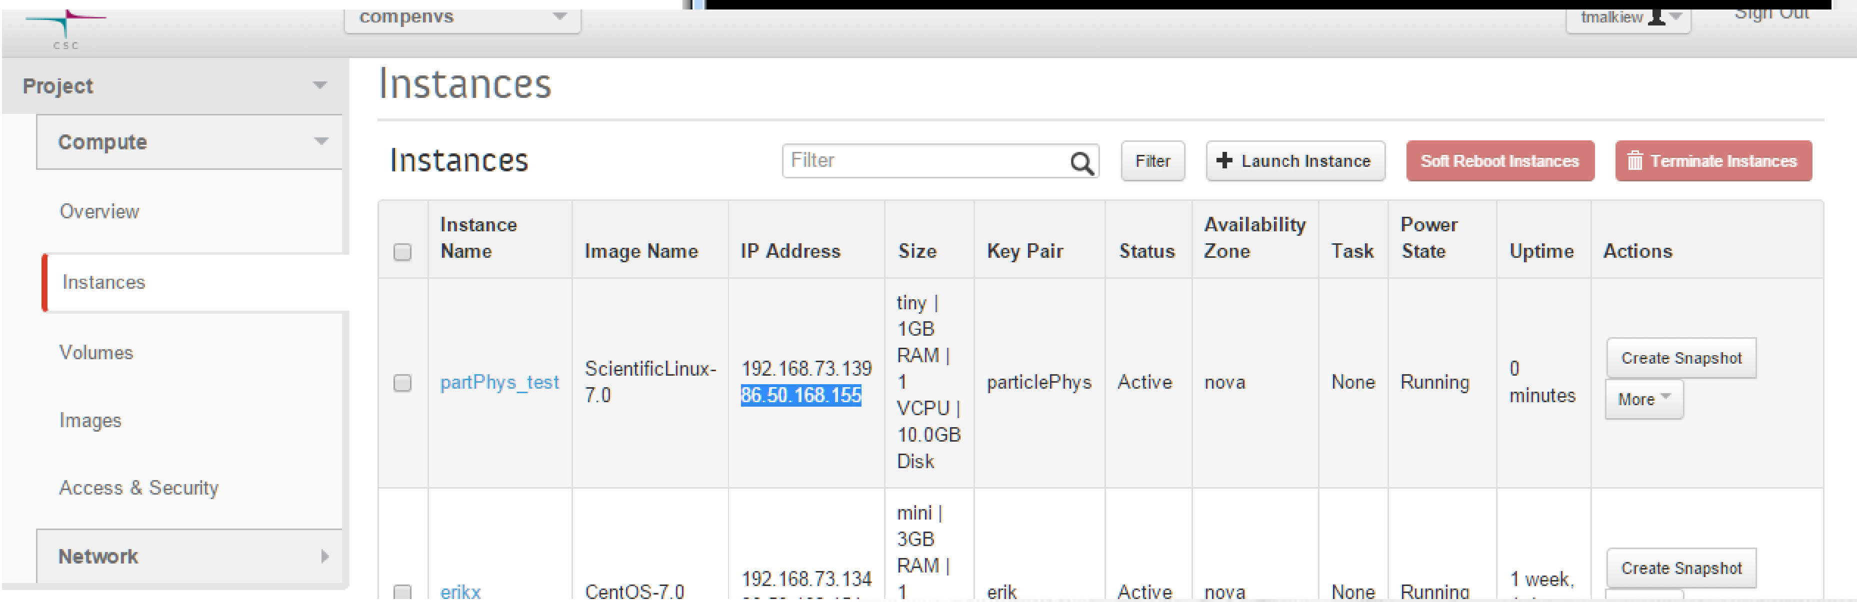
compenvs
tmalkiew
Project
Compute
Overview
Instances
Volumes
Images
Access & Security
Network
Instances
Instances
Filter
Filter
Launch Instance
Soft Reboot Instances
Terminate Instances
Instance Name
Image Name
IP Address
Size
Key Pair
Status
Availability Zone
Task
Power State
Uptime
Actions
partPhys_test
ScientificLinux-7.0
192.168.73.139
86.50.168.155
tiny | 1GB RAM | 1 VCPU | 10.0GB Disk
particlePhys
Active
nova
None
Running
0 minutes
Create Snapshot
More
erikx
CentOS-7.0
192.168.73.134
mini | 3GB RAM | 1
erik
Active
nova
None
Running
1 week,
Create Snapshot

Euclid project - cPouta use case 1

- cPouta being tested now for the Euclid Science Data Center challenges
- The European wide consortium will have a **common software environment**

  → virtual machine with administrator rights is a perfect match
- **Storage capacity** also essential

- The current resources for the project were applied by UH via standard CSC resource allocation process
- Later project is due to have **dedicated hardware** to be integrated to CSC 's cloud

CMS - cPouta use case 2

- Own CMS software stack
- HIP has been running a CMS Tier-2 site since 2008

- The virtualization decouples the CERN OS from the cloud host OS
- Allows to decommission the HIP owned Jade cluster hosted by CSC

- The Linux containers could be interesting to decrease the overhead of virtualized I/O
- Found challenges: The rapid development of OpenStack

CSC Finland – pioneer in Scientific Cloud Computing

- Ongoing/coming projects
  - DO-29401 (10 k cores – too early)
  - MS-4143
- Why to use CSC cloud services?
  - High-end compute HW, storage, networks
  - Based on open standards
  - Security is a high priority
  - Energy-efficient production
  - Specialist support & consulting
  - Services for all usage levels

Thank you for your attention!

Backup

Getting started

Academic use (OKM) free

- Apply for an account at CSC, then request Pouta access
  - https://research.csc.fi/pouta-application
  - Initial quota for testing & minor work
  - More quota can be applied for

Other use at cost

- Paid packages
  - Annual base package
  - Compute and storage quota options
- Send mail to contact@csc.fi for details

cPouta Use Cases

- Running scientific applications
  - Custom Operating System
    - Ubuntu, Debian, Scientific Linux, Windows Server 2012
  - Root/admin access needed for installation
- Building custom services
  - Web servers, file servers etc.
  - Software Defined Infrastructure (DevOps)
    - Rapidly deploying dev/test/prod environments
- Coursework