Medium-term evolution: Compute

session introduction

WLCG Collaboration Workshop
February 1, 2016

Oxana Smirnova
Lund University/NeIC

norden

NordForsk

neic Nordic e-Infrastructure
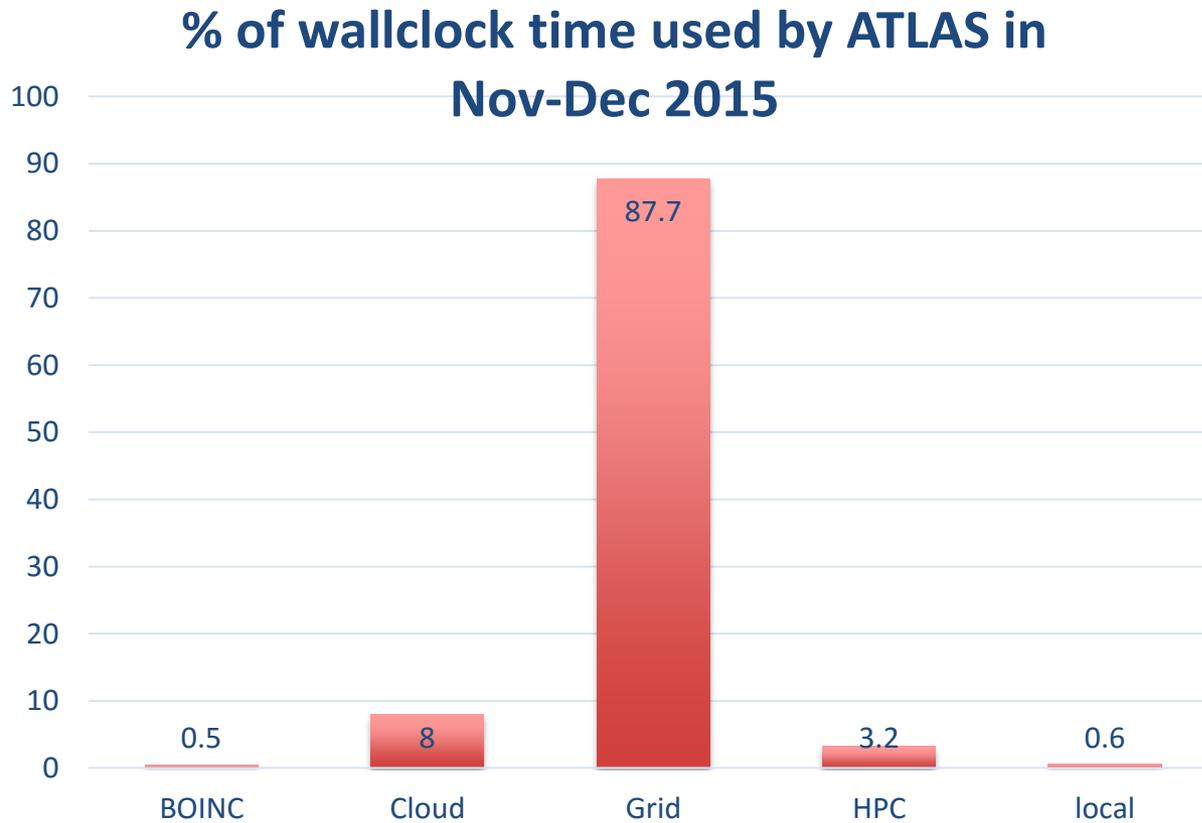Collaboration

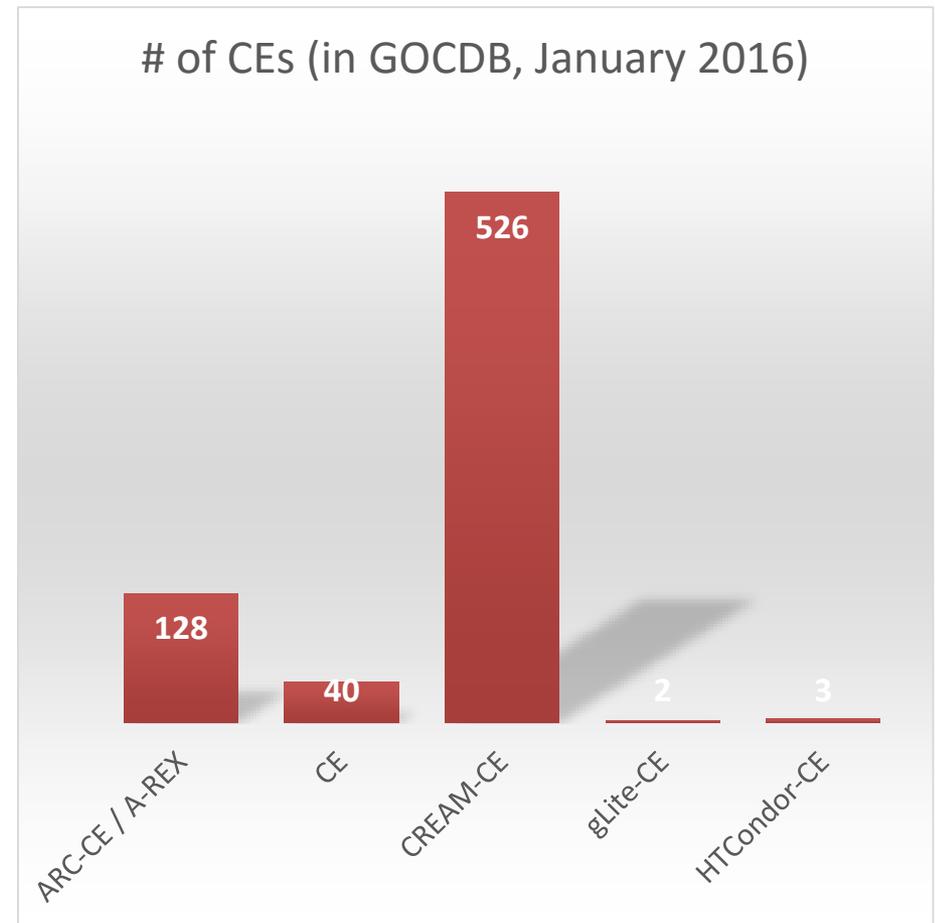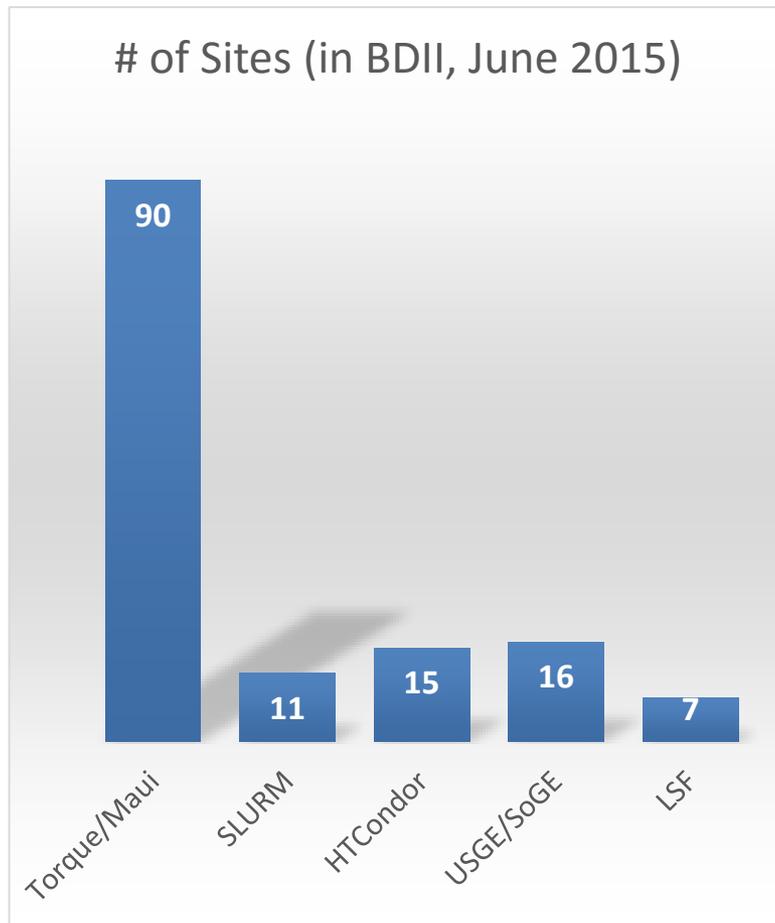LUND UNIVERSITY
Faculty of Science

Evolution of WLCG computing in medium-term

- Initial input from WLCG:
  - Simple access to any resources: grid, cloud, HPC, volunteer, etc
    - what is minimal need?
  - Need for "CE"
  - Need for batch systems?
  - Volunteer computing service as mechanism for "small" sites
  - Cloud provisioning: interfaces and how to provision; commercial clouds as extensions of sites, standalone?
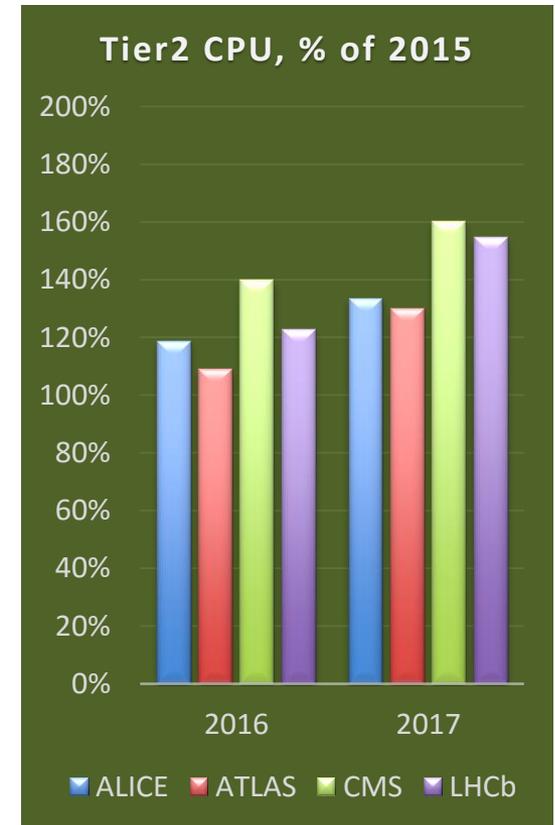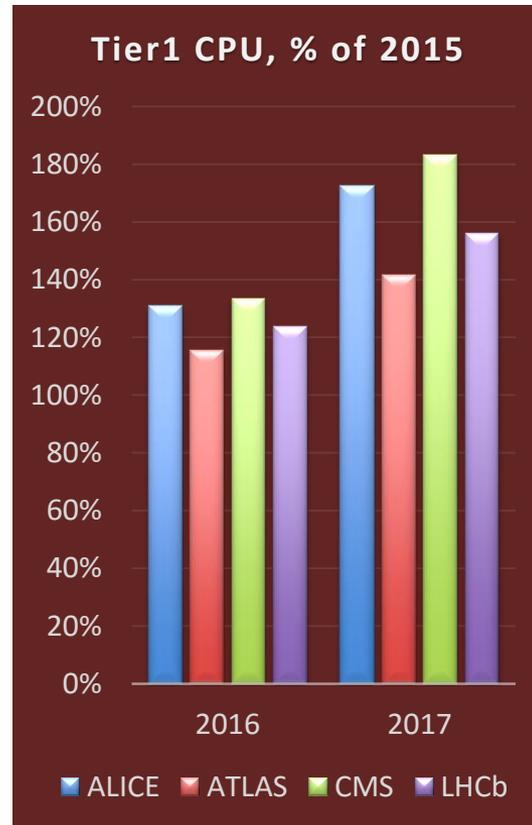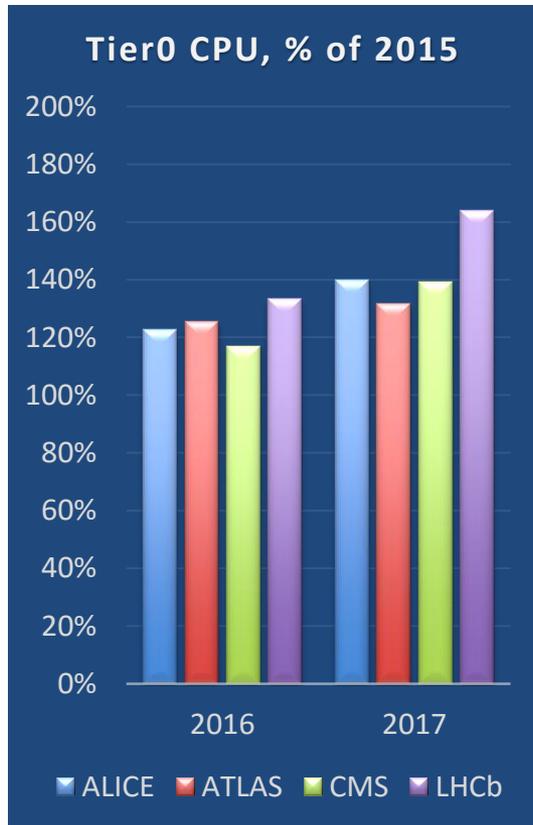
Different resources (ATLAS view)

% of wallclock time used by ATLAS in Nov-Dec 2015

Resource	%
BOINC	0.5
Cloud	8
Grid	87.7
HPC	3.2
local	0.6

Clusters and CEs

# of Sites (in BDII, June 2015)

Category	Value
Torque/Maui	90
SLURM	11
HTCondor	15
USGE/SoGE	16
LSF	7

# of CEs (in GOCDB, January 2016)

Category	Value
ARC-CE / A-REX	128
CE	40
CREAM-CE	526
gLite-CE	2
HTCondor-CE	3

* ~15% of CEs are not in production

Experiments' requirements

* 2017 requirements are preliminary

Session program

- Two key talks:
  - *"Exploiting HPC for HEP workloads"*, by Rod Walker
  - *"Plugging Cloud and VM-based resources into WLCG"*, by Andrew McNab
- Discussion panel
  - How much does WLCG need elastic resources?
    - or do we just need an extra capacity on a permanent basis?
  - Can common HPC/HTC resources totally replace dedicated HEP ones?
    - or is it politically easier to buy own hardware?
  - Why HPC resources are so HEP-unfriendly?
    - no outbound connectivity, diskless nodes, ban on services…
  - Is a cloud-HPC cross-over feasible?
    - running container-based VMs on HPC