



Track 3

Distributed

Computing



Stats

- 44 presentations
- 34 posters
- 7 sessions
- 4 conveners
 - Julia Andreeva (CERN)
 - David Cameron
(University of Oslo)
 - Hannah Short (CERN)
 - I. Ueda (KEK)
- 1 imposing venue



Topics

Supercomputing

Facilities

Accounting,
Testing,
Monitoring & IS

Performance
Optimisation

Security &
Identity

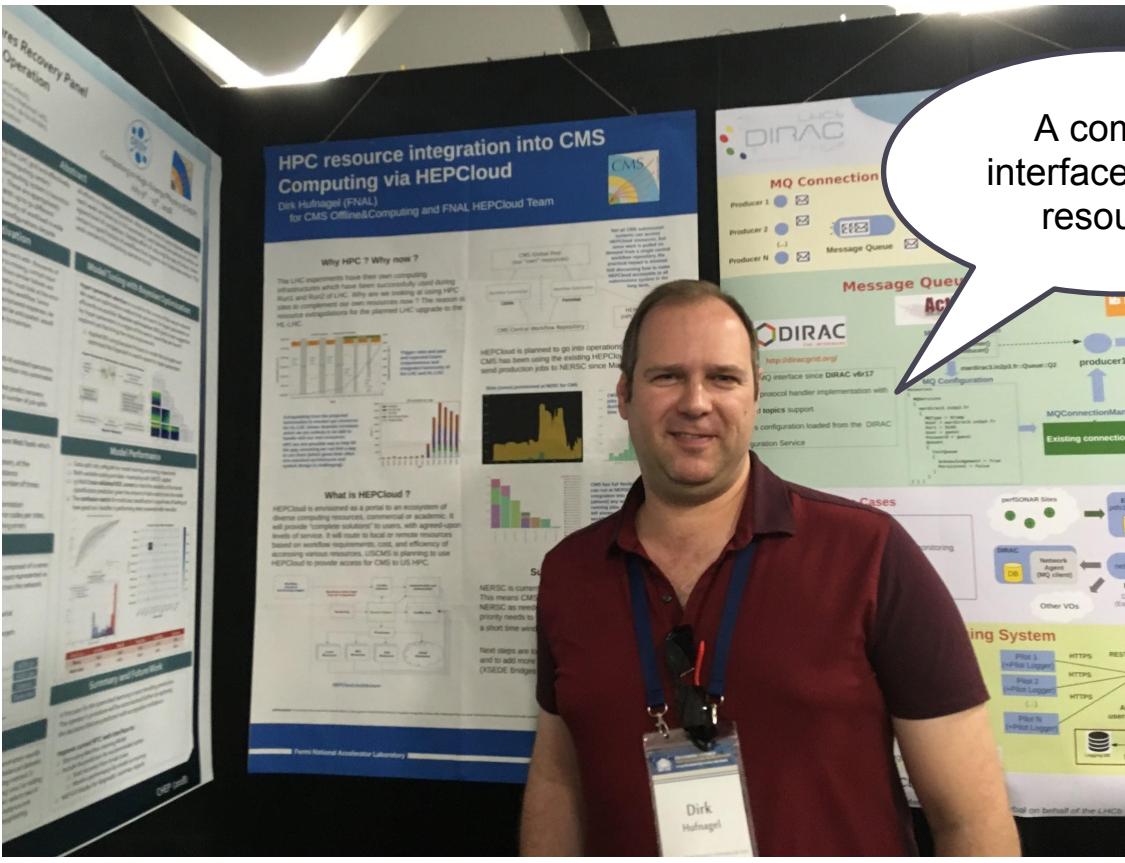
Experiment
Frameworks

Operational
Experiences

Computing
Models & Future
Views

Supercomputing

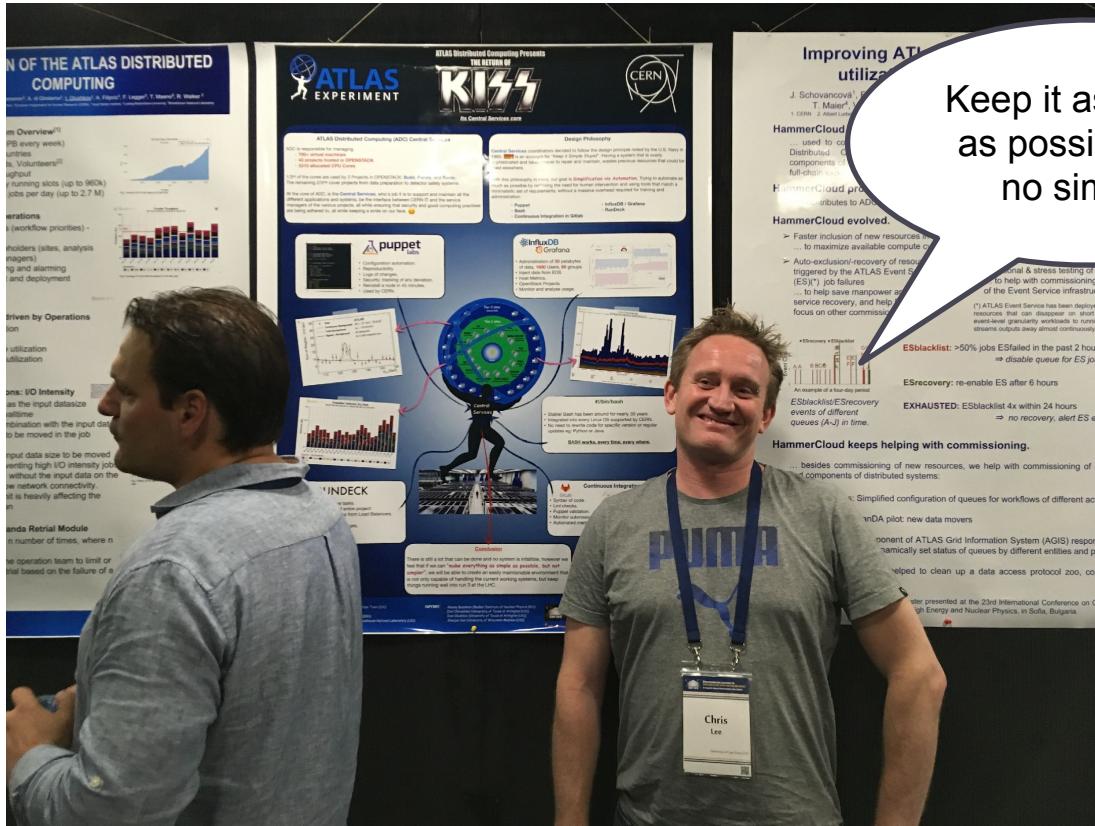
- Funding is flat but demand increasing. Could HPC be an answer to the challenges faced by HTC?
- Classical supercomputing workflows != Classical HEP workflows
 - Designed for single job that runs on 10s of thousands of CPUs
 - Environment differs significantly
- Supercomputing centres are beginning to support HEP but each one is different
- **There are many production use cases already**



HPC resource integration into CMS Computing via HEPCloud

Facilities

- Integration of **multiple** resource categories; HPC, Volunteer, Grid, GPU Clusters...
- Significant work on simulation of workflow distribution
- Shift in volunteer computing **away from true @Home** since fewer individuals have desktop PCs (and want to save battery!) Good backfilling opportunity.
- Successful efforts to extend HEP workflows and middleware to **other research communities**
 - LZ managed to plug a Tier 1 on top of a WLCG Tier 2 and it just worked, “distributed computing is awesome”



Accounting, Testing , Monitoring and Information System

- Data volumes continue to grow. Scalable, durable components are increasingly important, **modular architectures emerging as the correct approach**
 - Accounting - For dataset usage as well as resources
 - Testing - Take the burden off the experiments, test virtual resources pre-commission with HammerCloud
 - Monitoring - Experiments building up successful monitoring systems using common technology stack (message queues, Elasticsearch, Kibana, Grafana)
- CRIC proposes a **unified information system** for WLCG experiments. Will a combined approach work this time?



The GLUEX experiment can use HTCondor for file transfer (even if data rates doubles!)

Limits of the HTCondor transfer system

Performance Optimisation

- Lightweight monitoring for every job + powerful analytics = **identify bottlenecks and improve workflows**
- Exploit variability in jobs to efficiently fill slots.
Interruptible MC jobs provide an interesting **backfill** opportunity
- Using commercial SaaS proving cost effective... and encouraging implementers to build **efficient systems!**

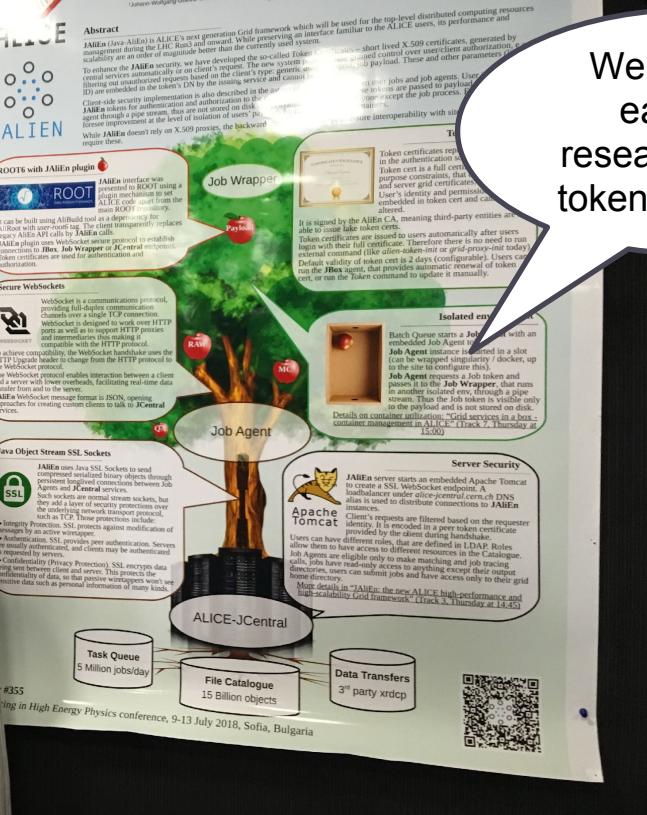
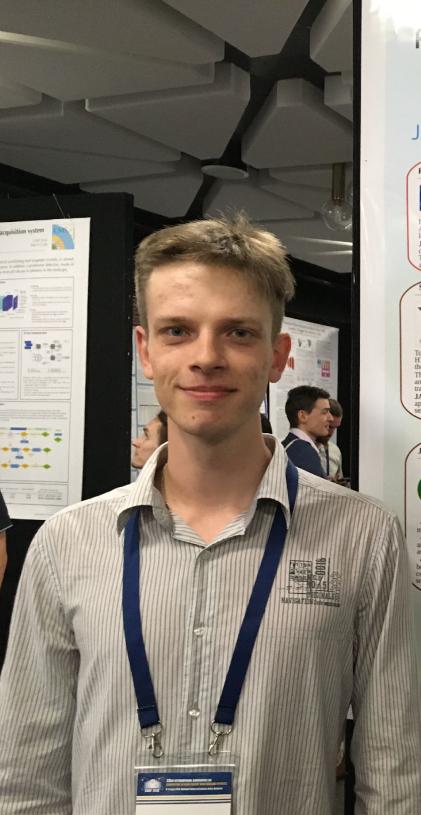


24PB for ALICE
during year 1 of
run 3 (place your
bets now!)

A simulation tool for ALICE storage and computing resource usage

Security & Identity

- WLCG Security Operations Centre WG hosting **SOC workshops** for sites to implement their own Intrusion Detection Systems for the benefit of all
- Research communities want to use Federated Identity Management to provide a good user experience to Researchers – but need the help of the wider stakeholder groups! **Reuse of AAI is key.** EOSC-Hub is working towards creating AAI solutions to unify Infrastructure access
- Macaroons as bearer tokens can put the user at the centre of authorising HEP workflows (BoF). Storage systems are enabling token based authorisation



We make life
easier for
researchers with
token certificates

The Security model of the ALICE next generation Grid framework

Experiment Frameworks & Operational Experiences

- Including: IceCube, ATLAS, XENON1T, LHCb, CMS, LSST, JUNO, COMPASS, CTA
- Each experiment has their own ways of working
 - Certain tools initially developed for a single LHC experiment are proving popular amongst multiple communities and outside LHC scope (e.g. Rucio, DIRAC, PanDA, Harvester)
 - Seeing a heartening increase in **reuse**
 - New experiments tend to do their own thing until hitting issues of scale when they look to LHC tools, how can they make the best choices?
- Trying to combine LHC workflows (in line with the HSF recommendations) will prove challenging(!)

Run2 jobs
overloaded the
conditions database
- we solved that!



Understanding the evolution of conditions data access through Frontier for ATLAS

Future Views

- Focused on **LHC for Run3**
- Performance Evaluation and Cost Modeling prediction is difficult, many parameters to be taken into account, new models are needed
 - Working Group in collaboration with HSF and WLCG
 - Defining common framework for modelling LHC experiments computing requirements and site cost estimation models
 - Many small improvements can make a big difference
- Simplified administrative workflows; Programmatic modeling instead of spreadsheets (CMS but generic)

Future Views

- Doing LHC computing ***even* better** this time around
 - Smarter file formats, e.g. reducing MC output by x20 (LHCb)
 - Flexible resource incorporation (ATLAS)
 - Containers playing a major role (CMS)
 - Efficient database design (ALICE)
 - Data movement optimisation being considered by DOMA (Data Organisation, Management and Access) Working Group
- Experiments like IceCube for which workloads use of GPU is ideal enable GPUs on the GRID

Track 3 Summary

- Funding flat but demand is increasing; efficient use of multiple forms of distributed computing is critical
 - Grid, GPUs, Supercomputing, Volunteer Computing...
- There is a focus on sharing LHC solutions and making them available to other communities
- Experiments are updating their technology stacks with modular, scalable software (often industry standard) that aims to achieve the optimisation needed for future computing demands



Saints Cyril and Methodius,
credited with the
cyrillic language

Traditional
Fire Walker

Lyudmila Zhivkova, who
proposed the idea of a
conference venue in
Sofia and is quoted as
saying “think of me as
fire”

The Fire, mural by Hristo Stefanov

THANKS!