Summary of Track 7 Clouds & Virtualization

Andrew McNab
University of Manchester

Claudio Grandi
INFN Bologna

Federico Stagni
CERN

Jeff Templon
NIKHEF

The word “cloud” in CHEP abstracts

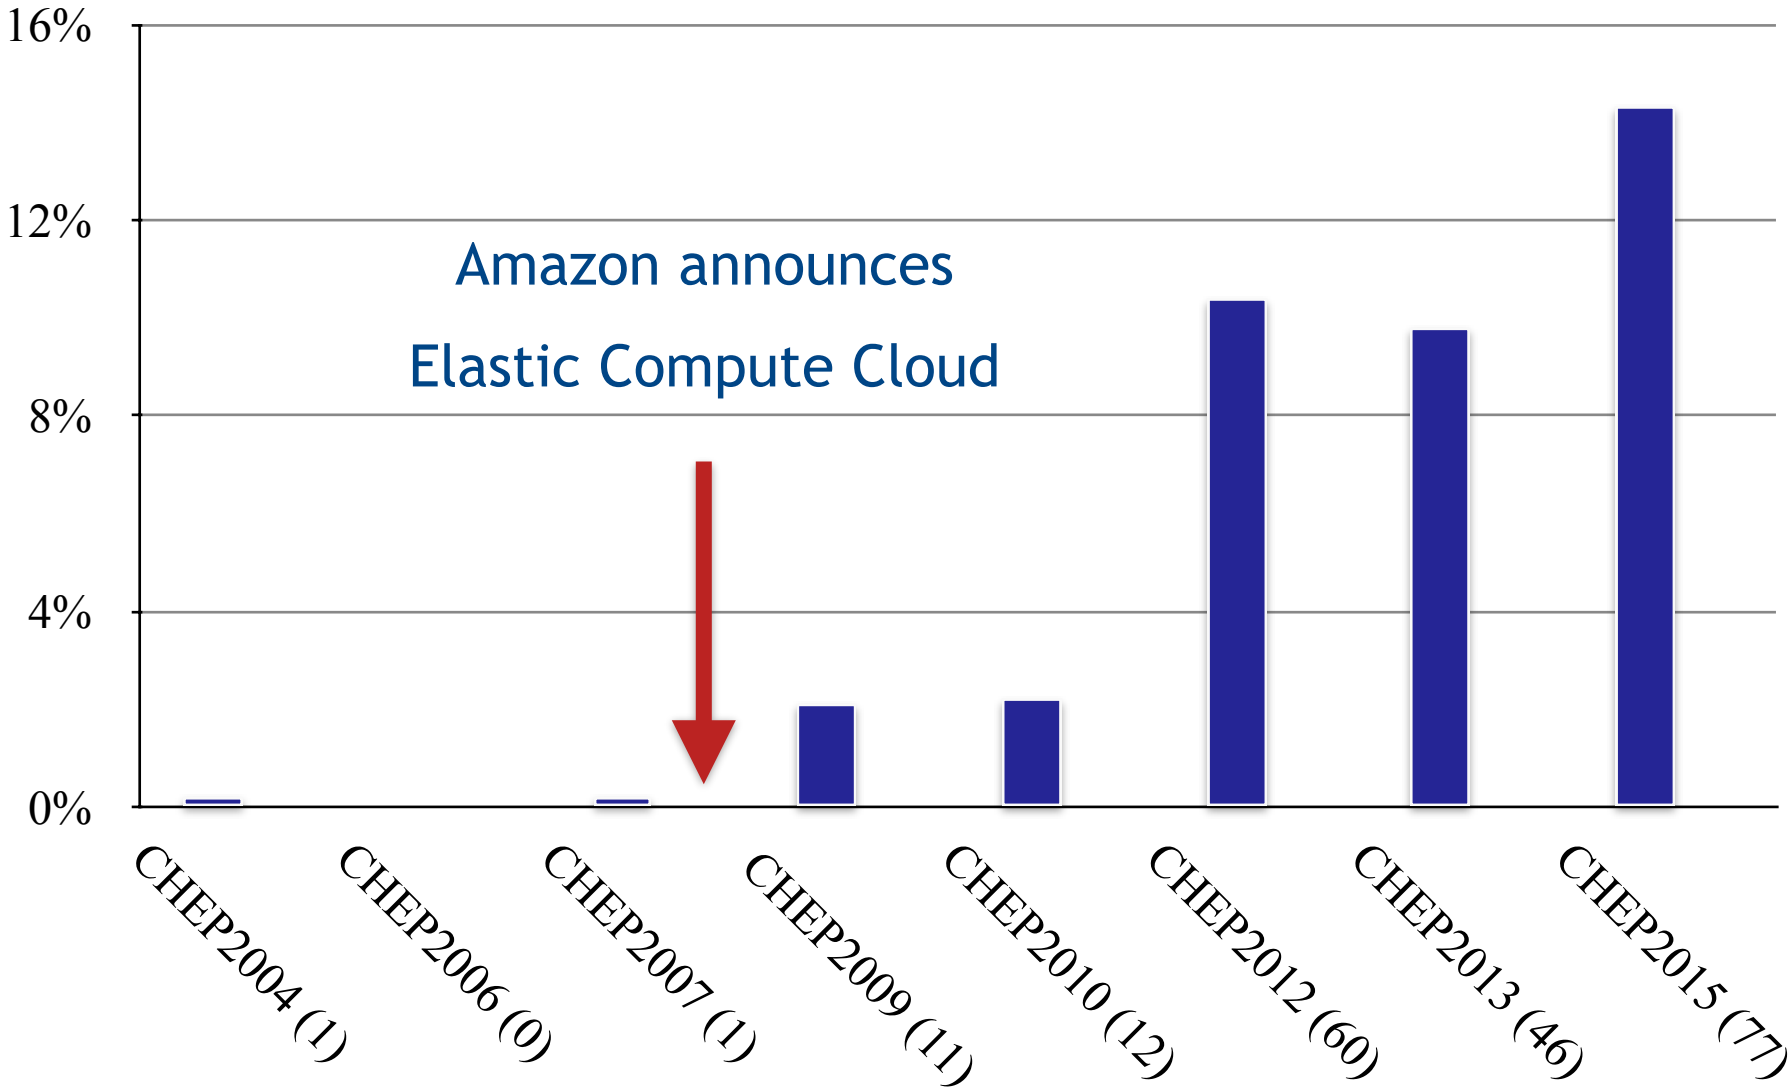

Gartner Hype Cycle

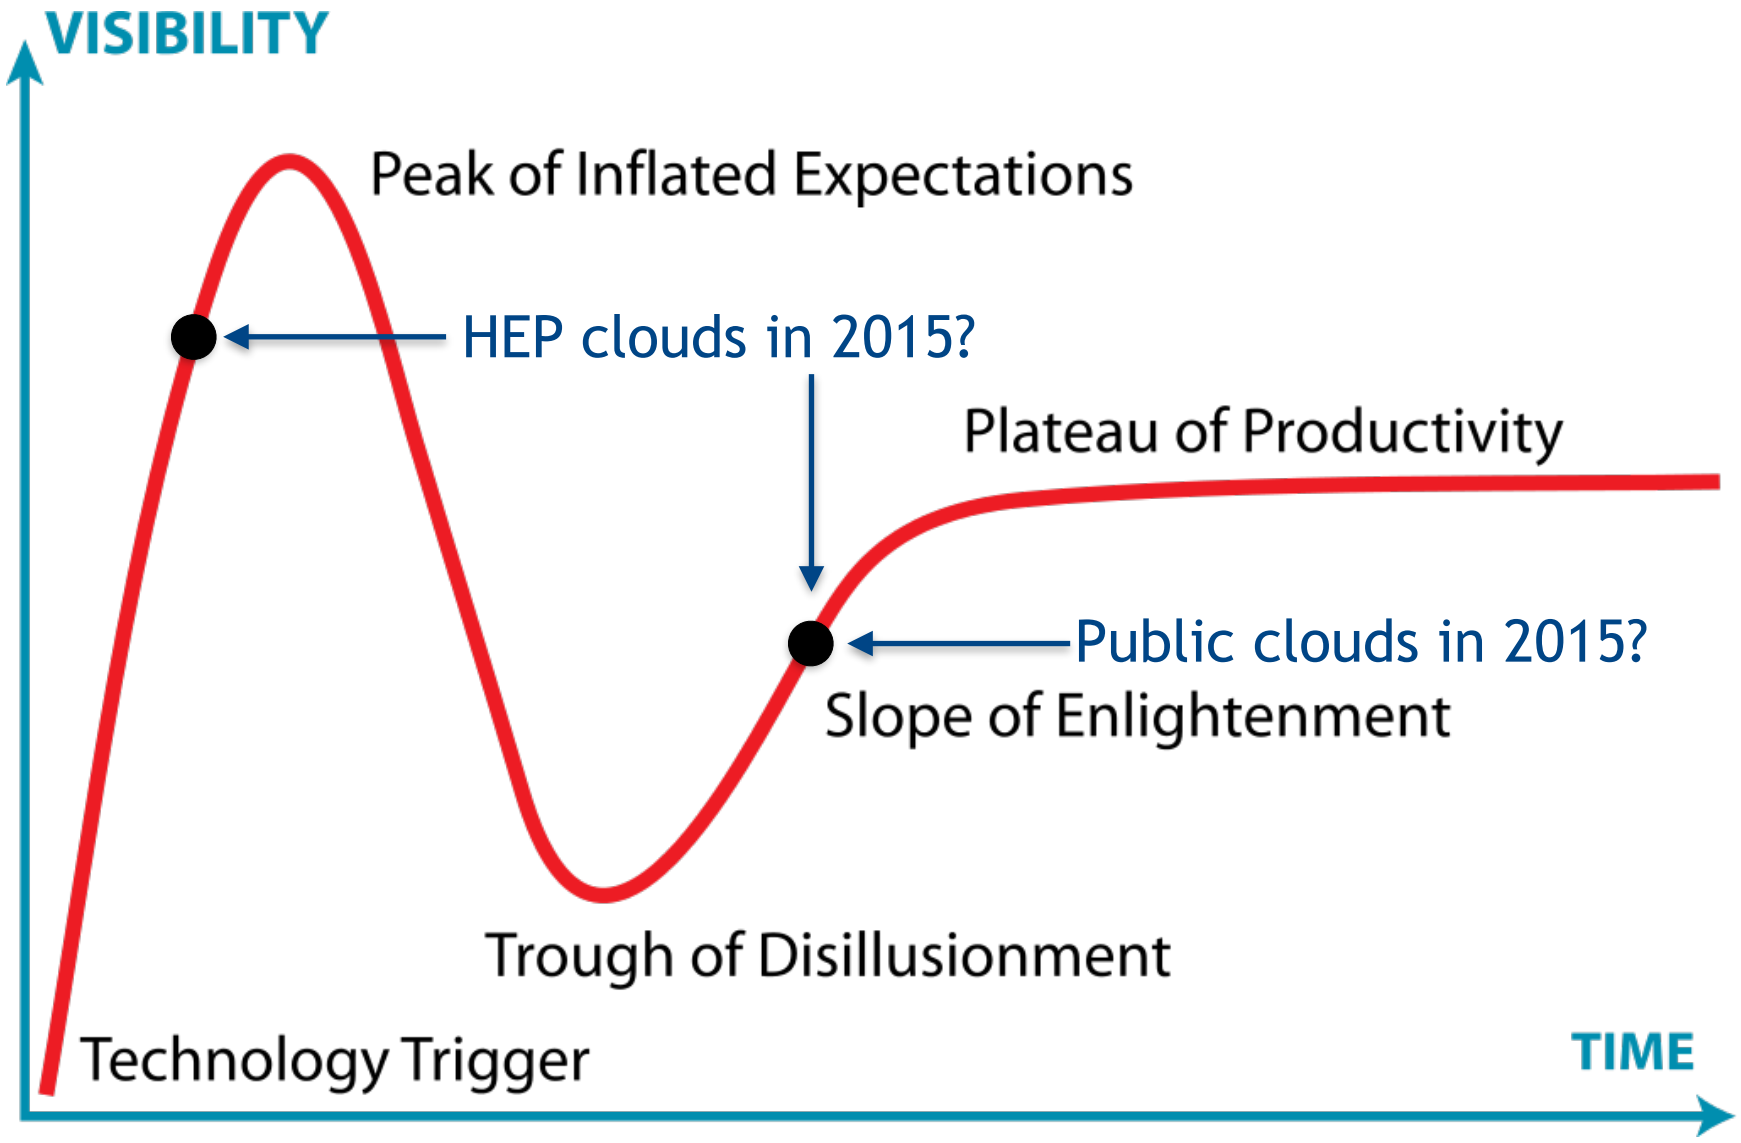

Track 7: Clouds and Virtualization

- First time that there's been a dedicated track on this
  - Previously (2010,2012) there had been "… with Grids and Clouds" tracks
- We collected abstracts submitted with "Cloud" (81) or "Virtualization" (32) keywords
  - Included Containers as a type of Virtualization
- Negotiation with other tracks on borderline topics
- Divided abstracts into talks (30) and posters (31)
- Assigned talks to one of four sessions:
  - site operations; volunteer computing and storage; experiments; technologies

How did it go?

- The whole process went very smoothly
- I'd like to thank Claudio, Federico & Jeff for being a pleasure to work with
- Everyone took the lead during the process at one time or another

- The sessions were very well attended:
  - People were standing/sitting on the floor at some point during all of the sessions, in a room with 70 chairs
- The sessions were lively and interesting, and a lot of new ideas exchanged
- Had some last minute cancellations
  - So 15 minute slots for extra questions / ad-hoc panel discussions
  - Very useful in practice!

Themes

- Emphasis on running jobs
  - Rather than coverage of virtualization as a platform for services
- VM-based jobs systems are used for routine production work
  - CloudScheduler, GlideinWMS, and the 3 Vacuum-based platforms seem to be biggest numerically
- OpenStack prominent at sites but other IaaS platforms are common too
- Commercial cloud use still almost all relies on (large scale) donated resources from providers
- Volunteer computing with BOINC is rapidly emerging as very large potential source of capacity
- Many talks mentioned that they could or would support containers as an alternative to full-blown VMs
- Container-specific talks mostly emphasised packaged apps paradigm

Session 1: Monday

Site operations

Scaling the CERN OpenStack cloud

Stefano Zilli
On behalf of CERN Cloud Infrastructure Team

Stefano explained how OpenStack cells are allowing the CERN private cloud to scale ready for Run 2

CERN Private Cloud - Numbers

- Based on OpenStack Juno
- Spans between 2 datacentres
- 4700 hypervisors
  - 120000 cores
- 11000 VMs
- 1500 users
- 1800 projects

Scheduler

- Two schedulers
  - Cell scheduler
    - Top cell level
    - Decides the cell
  - Node scheduler
    - Compute cells
    - Decides the hypervisor

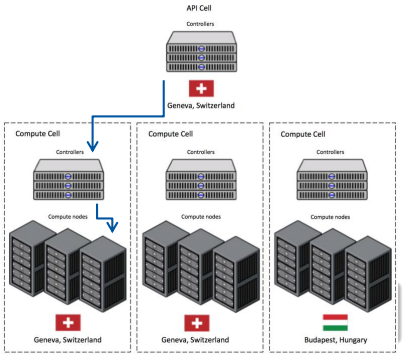

Benchmarking and accounting for the (private) cloud

Jerome Belleman, Daniel Pek, Ulrich Schwickerath,

Special thanks to the CERN Cloud team

Ulrich presented detailed work on how to classify and benchmark VMs, and then inject performance numbers into the accounting system.

Classification of worker nodes

Example: a6_8_1512h23_266
AMD based virtual machine

- SLC6
- 8 cores
- CPU-ID 1512h, see below
- CPU speed 2300 MHz
- Default memory speed 266

Remark : Details of the machine:

Memspeed = unknown
Cpuspeed = 2.3MHz
Cpuvendor = AMD
Cores = 8

CPUfamily = 21 =0x15h
CPUmodel = 1 = 0x1h
CPUstepping = 2 =0x2h

=> **CPUID = 1512h**

http://world.std.com/~swmcd/steven/tech/cpu.html

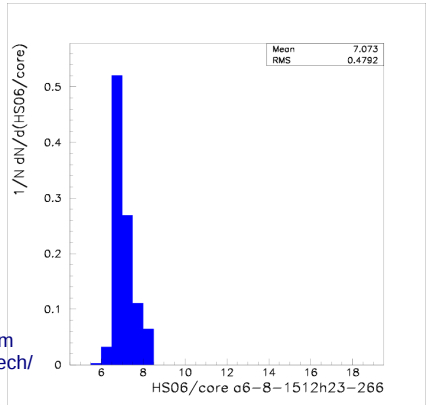

Cloud accounting: general case

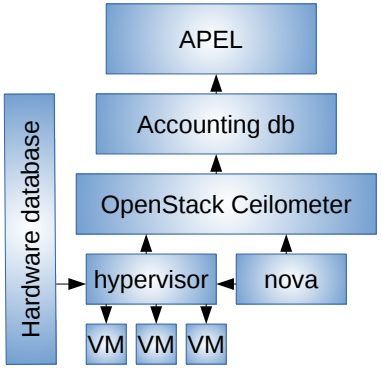

- Work in progress:
  - Inject performance info from the hypervisor to ceilometer while the VM is running
- Possible future work:
  - inject all information we need to do the classification as for the batch case
  - Unclear how to do this in a general case

Managing virtual machines with Vac and Vcycle

Andrew McNab
University of Manchester

Peter Love
Lancaster University

Ewan MacMahon
University of Oxford

I showed how Vac and Vcycle are using the Vacuum model to run production work for experiments in Pilot VMs at multiple sites.

Vac - the first Vacuum system

Infrastructure-as-a-Client (IaaC)

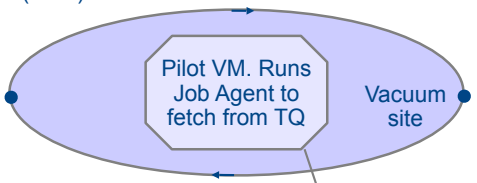

Since we have the pilot framework, we could do something really simple

Strip the system right down and have each physical host at the site create the VMs itself.

Instead of being created by the experiments, the virtual machines appear spontaneously "out of the vacuum" at sites.

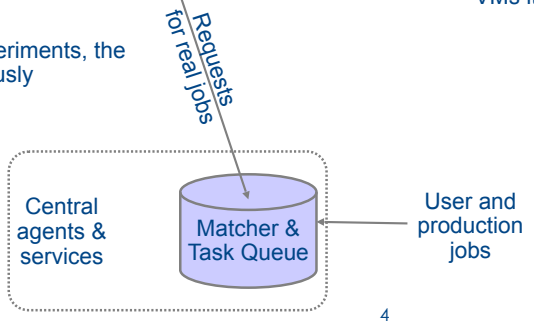

Use same VMs as with IaaS clouds

Deployment by site and experiment

		ATLAS	CMS	LHCb	GridPP DIRAC
Vac	Manchester	✔	✔	✔	✔
	Oxford	✔	✔	✔	✔
	Lancaster	✔		✔	✔
	Birmingham				✔
Vcycle	CERN (LHCb)			✔	
	CERN (Dev)	✔	✔	✔	✔
	Imperial	✔	✔	✔	✔
	CC-IN2P3			✔	
HTCondor Vacuum at RAL (STFC)		✔	✔	✔	✔

University of Glasgow

THE UNIVERSITY of EDINBURGH

Evaluation of containers as a virtualisation alternative for HEP workloads

Gareth Roy[1], Andrew Washbrook[2],
David Crooks[1], Gordon Stewart[1], Gang Qin[1], Samuel Skipsey[1], Dave Britton[1]
[1]University of Glasgow [2]University of Edinburgh

21st International Conference on Computing in High Energy Physics
13th April 2015

Andrew explained how Containers differ from Virtual Machines, and presented performance comparisons with (unoptimised) VMs, with benchmarks and with realistic HEP workloads.

Containerisation

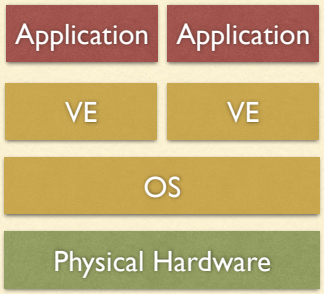

Linux Container

- Containerisation is a form of **OS level virtualisation**
- The Linux kernel hosts multiple partitioned user-land instances (Virtual Environments)
- Accomplished through separate *namespaces* for filesystem mounts, network, processes and users
- Backing storage can be Copy-on-Write or a union filesystem (UnionFS/AUFS)

Containers Timeline

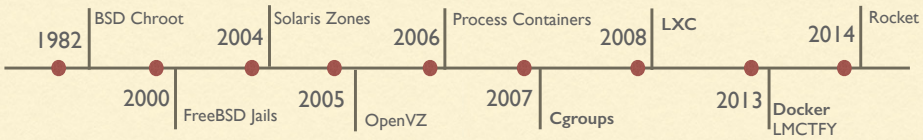

Event Generation and Reconstruction Results

Event Generation

- Measured total time taken to generate large sample of $Z\mu\mu$ events
- VMS lose only **5%** performance compared to bare metal and containers

Event Reconstruction

- CPU time/event averaged over 50 events
- VMs show **22.6%** performance drop for a single process, **18.8%** for 8 simultaneous processes

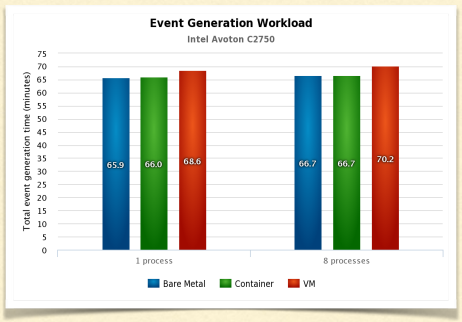

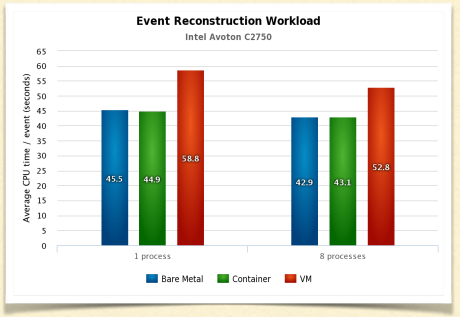

docker & HEP: containerization of applications for development, distribution and preservation

Sébastien Binet

LAL/IN2P3

2015-04-13

Ben showed us Sebastien's slides, explained how easy it is to create containers, and showed some benchmarking with Gaudi.

Docker: creating a customized container

- run `docker` interactively:

```
$ docker run -i -t ubuntu bash
root@bf72b1a06e6c:/# apt-get update
Reading package lists... Done

root@bf72b1a06e6c:/# apt-get install memcached
[...]
root@bf72b1a06e6c:/# exit
```

- commit the resulting container

```
$ docker commit `docker ps -q -l ` binet/memcached
ab59e4b14266
```

- run the image

```
$ docker run -d -p 11211 -u daemon binet/memcached memcached
ab59e4b14266
```


Benchmarks - III

Running `gaudirun.py GaudiExamples/TupleEx.py`

- `AFS`

```
56.87s user 14.26s system 66% cpu 1:46.50 total # kick AFS
57.62s user 13.07s system 99% cpu 1:11.17 total
57.69s user 13.46s system 99% cpu 1:11.58 total
57.93s user 13.26s system 99% cpu 1:11.66 total
```

- `Docker-RPMs`

```
55.93s user 12.34s system 98% cpu 1:09.54 total
55.43s user 12.88s system 98% cpu 1:09.12 total
55.54s user 12.16s system 98% cpu 1:08.83 total
55.39s user 11.60s system 98% cpu 1:07.81 total
```

- `Docker-CVMFs` (a `docker` container where `CVMFs` is configured and running)

```
55.53s user 14.01s system 88% cpu 1:18.75 total # kick CVMFs
54.95s user 12.83s system 97% cpu 1:09.36 total
55.42s user 12.86s system 98% cpu 1:09.35 total
55.42s user 13.01s system 98% cpu 1:09.63 total
```

Integrated Monitoring-as-a-service for Scientific Computing Cloud applications using the ElasticSearch ecosystem

Speaker: Sara Vallero

Bagnasco Stefano
Berzano Dario
Guarise Andrea
Lusso Stefano
Masera Massimo

The present work is partially funded under contract 20108T4XTM of Programmi di Ricerca Scientifica di Rilevante Interesse Nazionale (Italy).

INFN Istituto Nazionale di Fisica Nucleare

Sara described the INFN Torino private cloud, and showed how applications such as the Virtual Analysis Facility are being monitored.

The INFN Torino Private Cloud

WLCG Tier2
(ALICE, LHCb, biomed, CTA, Panda, Belle2)

Virtual Farms on-demand:
theory
Compass
medical imaging
JLab
…

Customised VMs:
ufsd
nuclear plant simulation
…

OpenNebula

1.3k cores

1.6k TB (gross)

1-10 Gbps LAN
10 Gbps WAN

DIRAC Tier2
(BESIII)
E

R&D

Virtual Analysis Facility
(ALICE)
E

Elastic Farm
(ALICE)
E

Applications monitoring: VAF dashboard

click to select entries for specific user only

dynamic time range (the selection is applied to all plots)

example of ES analysis on a complex string

path/to/file

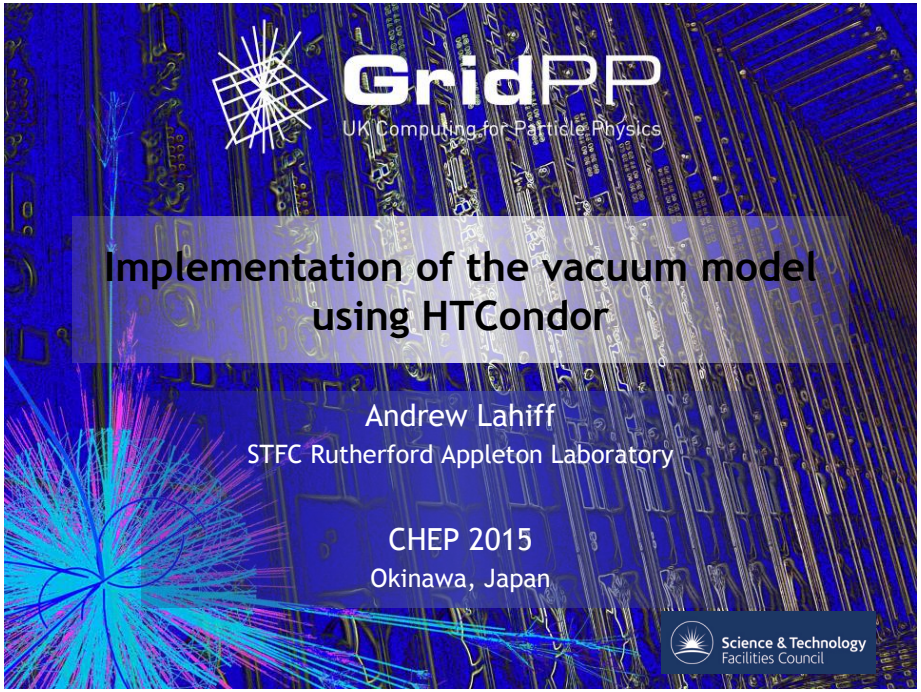

Andrew explained how he had applied the Vacuum model to make HTCondor create Pilot VMs in response to demand from experiments' workloads.

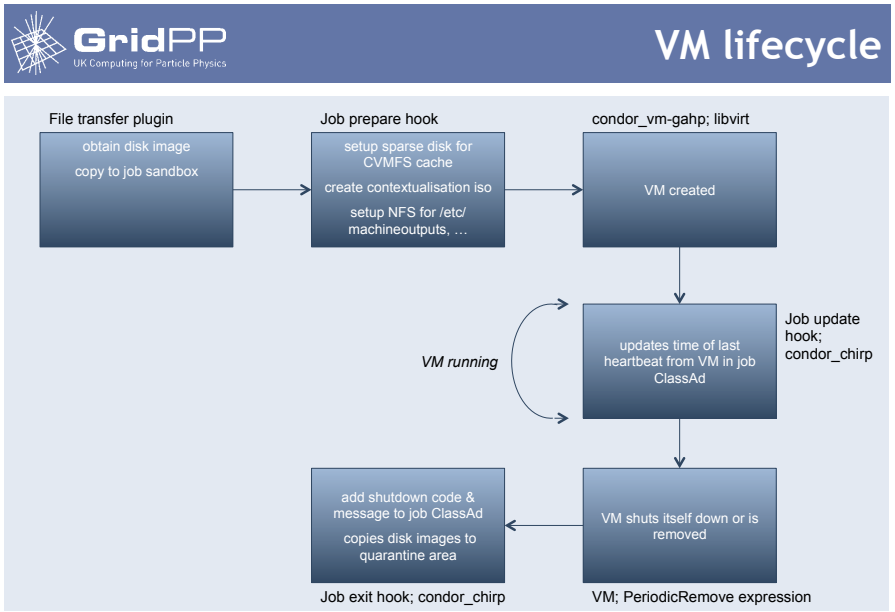

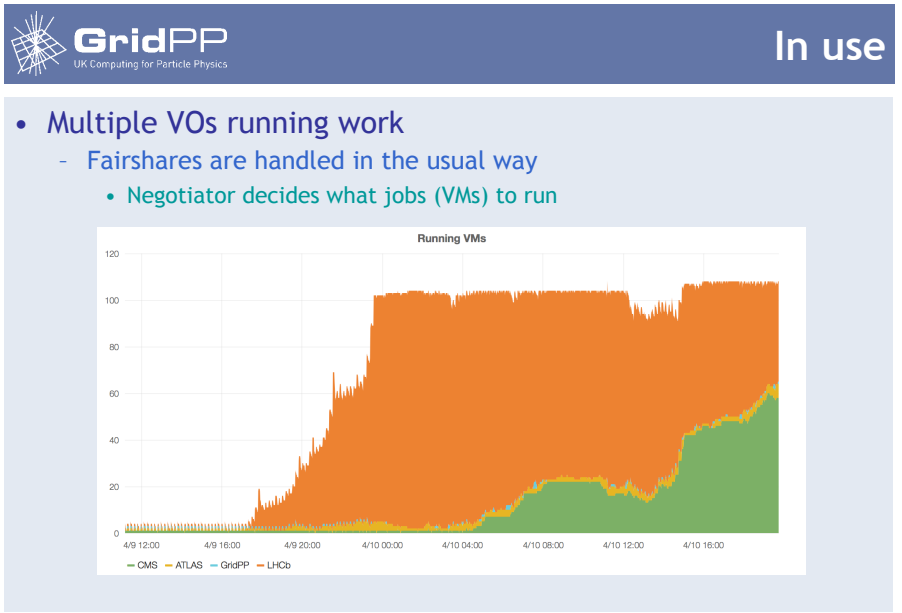

Session 2: Tuesday

Volunteer computing and storage

21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

Towards a production volunteer computing infrastructure for HEP

Nils Høimyr (CERN), Miguel Marquina (CERN)
Tomi Juhani Asp (University of Jyvaskyla, Finland),
Laurence Field (CERN), Alvaro Gonzalez (CERN), Pete Jones (CERN)

Presentation on behalf the CERN BOINC Service Team

Miguel explained how they are supporting volunteer computing with BOINC in HEP, enabled by virtualization.

BOINC – Volunteer view

- Download and run BOINC client
- Choose a project
- Provide email and password to the BOINC Manager
  (alternatively make a silent connection with a key from the BOINC client)
- Done, crunching can start!

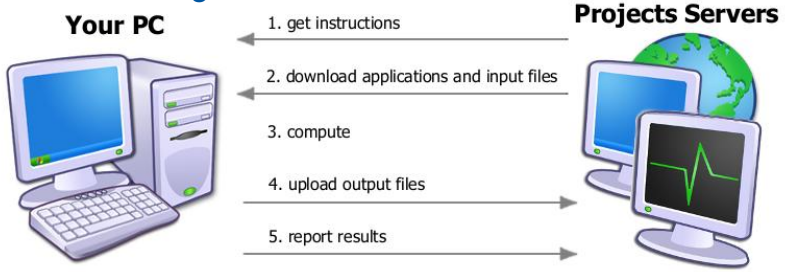

Virtualisation in BOINC - 2

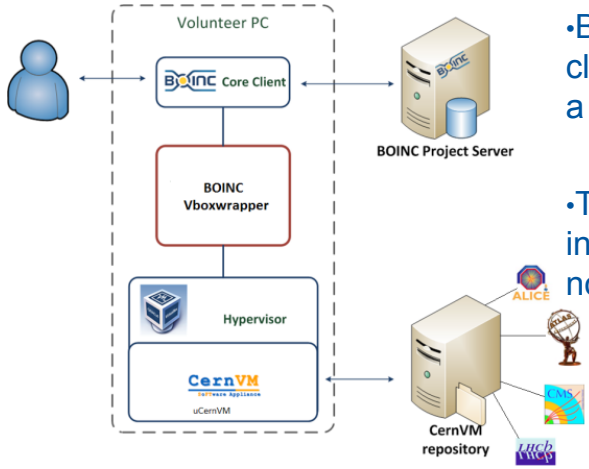

- BOINC distributes VMs to client machines along with a wrapper application
- The BOINC client installation for Windows now includes Virtual Box

UiO : Department of Physics
University of Oslo

David Cameron
Riccardo Bianchi
Claire Adam Bourdarios
Andrej Filipcic
Eric Lançon
Efrat Tal Hod
Wenjing Wu
on behalf of the ATLAS Collaboration

CHEP 15, Okinawa

ATLAS@Home: Harnessing Volunteer Computing for HEP

David showed how BOINC has become one of the largest contributors to ATLAS simulation in the first year of using it!

UiO : Department of Physics
University of Oslo

Basic ATLAS@Home Architecture

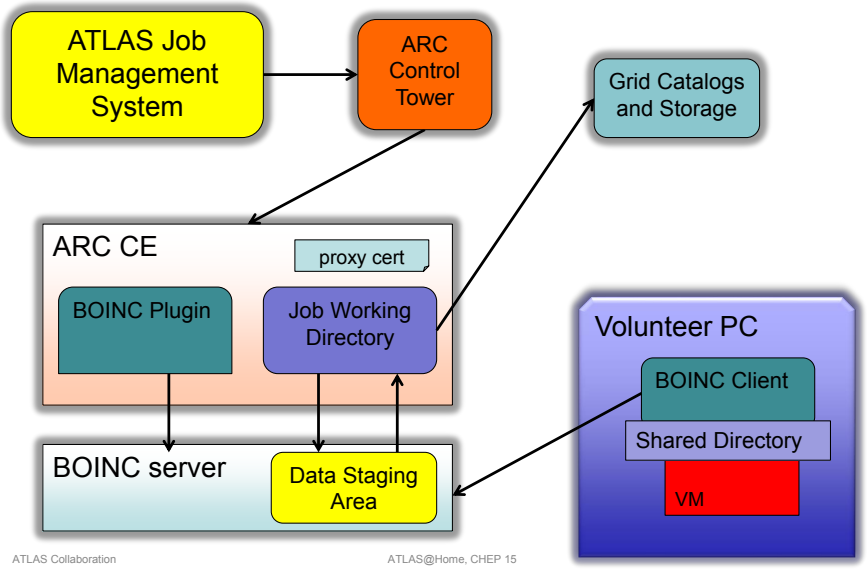

UiO : Department of Physics
University of Oslo

Scale of ATLAS@Home

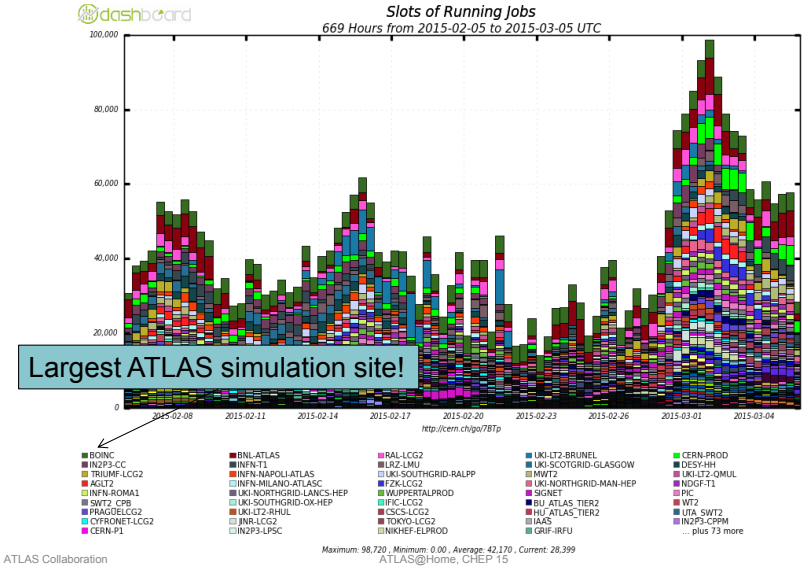

CMS@home: Enabling Volunteer Computing Usage for CMS

Laurence Field (CERN), Hendrik Borras (Heidelberg University), Daniele Spiga (CERN), Hassen Riahi (CERN)

Laurence described how use of the DataBridge is being prototyped by CMS to interface volunteers' VMs and their trusted services

Infrastructure

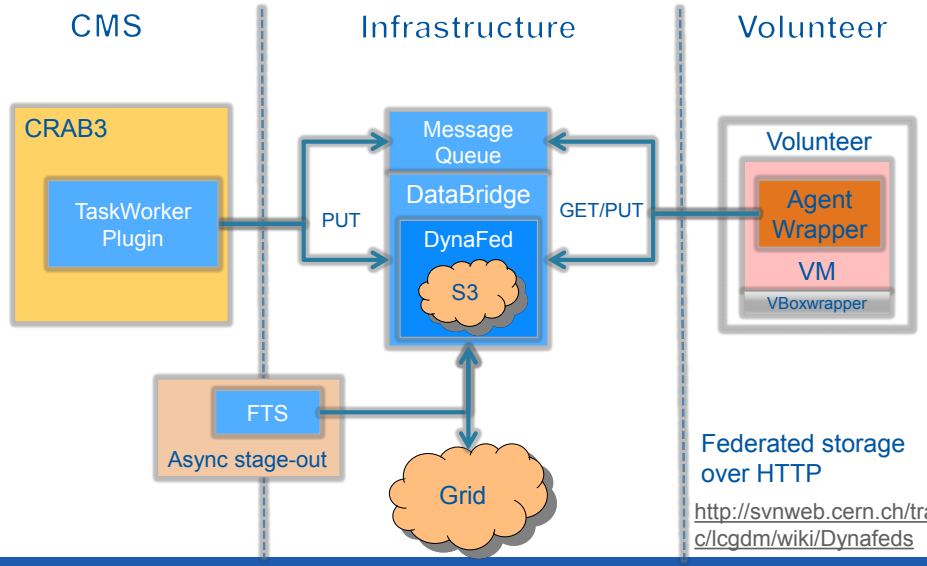

The DataBridge

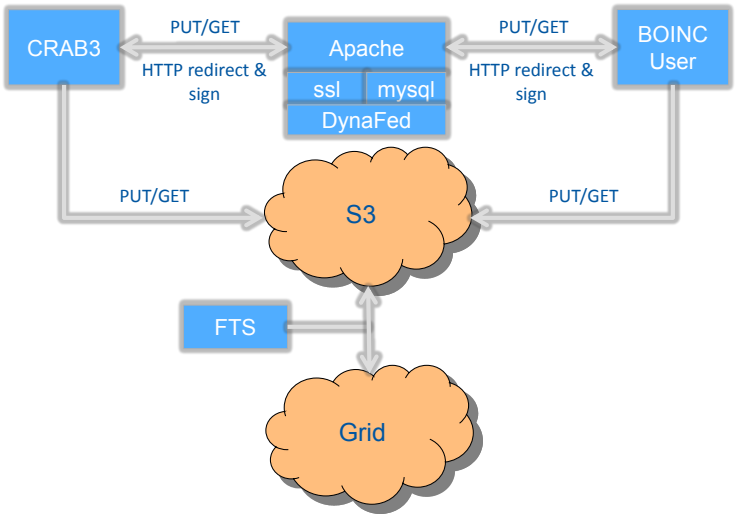

Karlsruhe Institute of Technology

Dynamic provisioning of local and remote compute resources with OpenStack

CHEP 2015 - Okinawa, Japan

Manuel Giffels, Thomas Hauth, Frank Polgart, Günter Quast | 14. April 2015

INSTITUTE OF EXPERIMENTAL NUCLEAR PHYSICS (IEKP)

KIT – University of the State of Baden-Wuerttemberg and National Laboratory of the Helmholtz Association

Thomas described setting up a multidisciplinary OpenStack HPC Cluster, with access for HEP by HTCondor-enabled VMs.

From HPC Cluster to HEP Worker Node

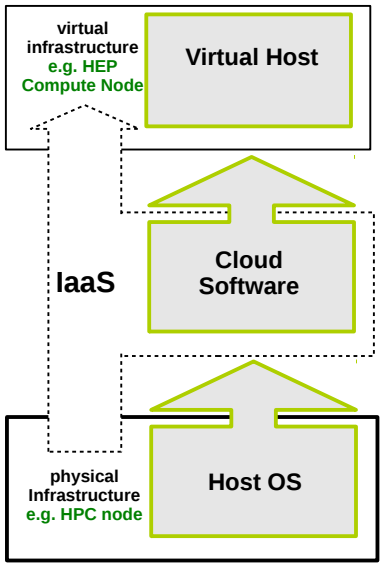

The Infrastructure-as-a-Service (IaaS) model

- Infrastructure (e.g. machines, network) is virtualized
- Decouples complexities of hardware maintenance and specific software setup
- The life cycle of this virtual infrastructure is managed by a Cloud system:
  - Virtual machine images are managed
  - The user can upload and start custom virtual machines
  - Storage blocks can be attached to these VMs

Topology of Local and Remote Sites

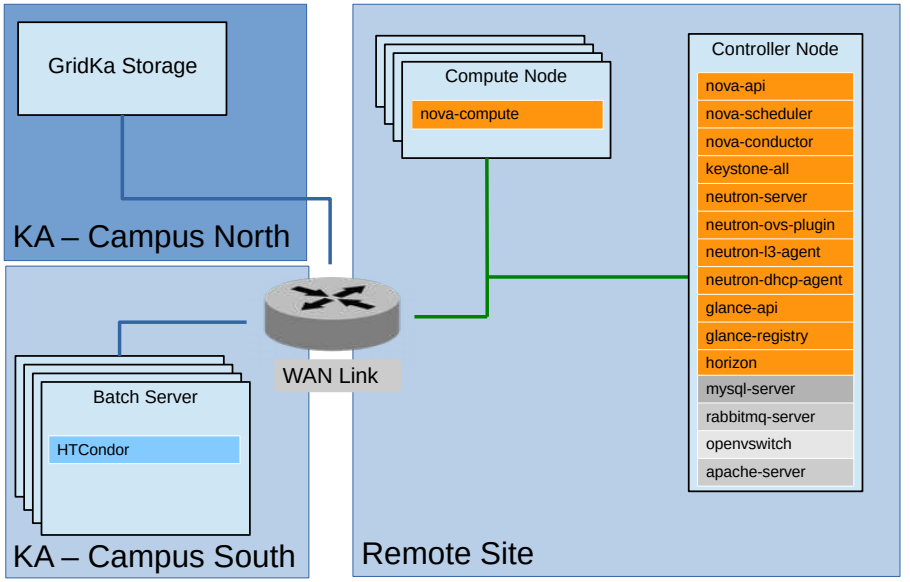

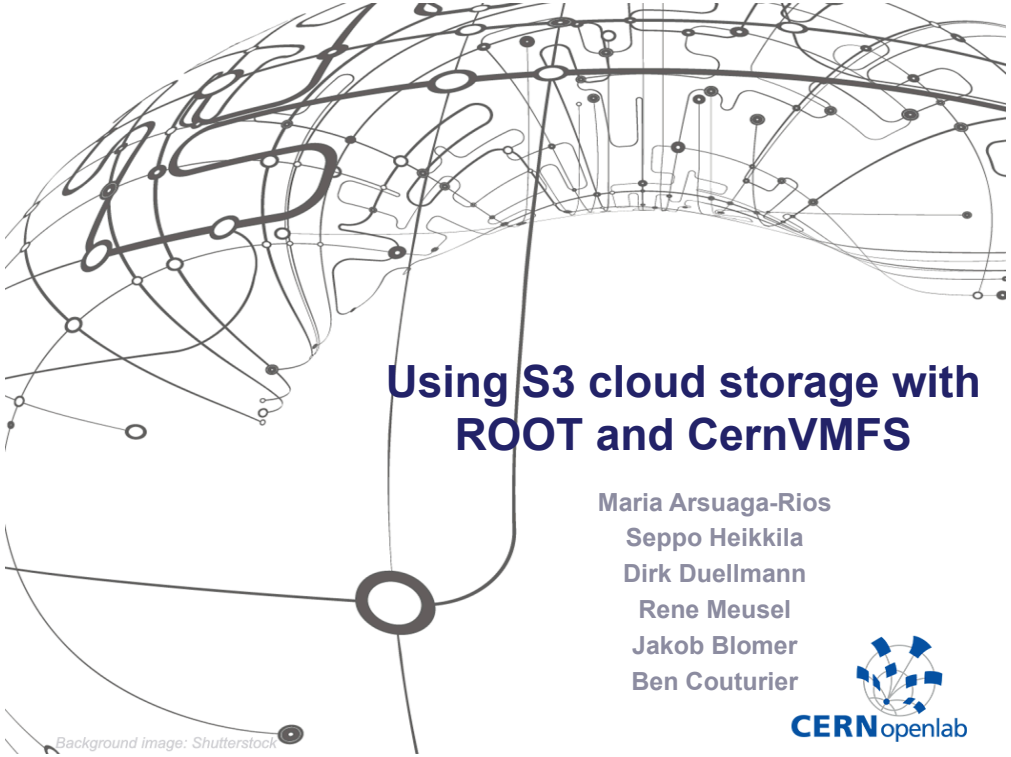

Using S3 cloud storage with ROOT and CernVMFS

Maria Arsuaga-Rios
Seppo Heikkila
Dirk Duellmann
Rene Meusel
Jakob Blomer
Ben Couturier

CERN openlab

Background image: Shutterstock

Maria presented ROOT file access benchmarks for the S3 service; and how CernVM-FS on S3 is being used, for example for LHCb nightly builds.

CERN openlab

MULTI-CLIENT ROOT BENCHMARKS 64 Buckets

Old UDS

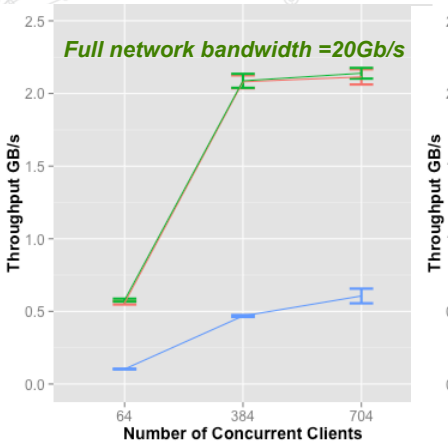

New UDS

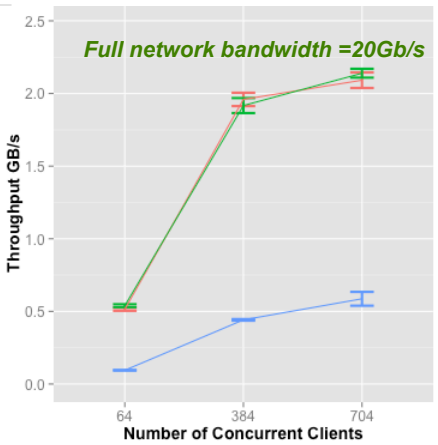

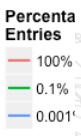

CERN openlab

REAL CernVMFS APPLICATION S3 Backend

- **New features to CernVMFS (release 2.1.20)**
  - S3 compatible storage
  - Supports multiple buckets and accounts
- **LHCb nightly builds test case**
  - LHCb software (binary and source files)
  - Releases done daily for over 3 months
  - ~1 million files in release (~150k new files)
  - Completed successfully with UDSs and Ceph

Background image: Shutterstock

New adventures in storage: cloud storage and CDMI

Paul Millar
On behalf of the dCache team.

CHEP 2015: Spring Okinawa

Fermilab NDGF Nordic DataGrid Facility DESY HELMHOLTZ ASSOCIATION LSDMA

Paul outlined the dCache team's plans to support the CDMI protocol, and how CDMI addresses many things on HEP's wish-list for storage

dCache.org

What is dCache plan?

- Adding **initial CDMI support** this summer,
  - filesystem-like access
  - Based on collaboration with HTW Berlin,
- Next steps:
  - Object store,
  - Storing metadata,
  - Querying metadata.

dCache.org

CDMI supports

- Writing data with different cost/retention guarantees: **tick**.
- Managing latency guarantees: **tick**.
- Third-party copy: **tick**.
- Notification when missing files: **partial-tick** (logging).
- Notification when files deleted: **partial-tick** (logging).
- Discover file popularity: **partial-tick**.
- Storage accounting: **tick**.
- Access monitoring: **partial-tick** (logging open/close).

Enabling opportunistic resources for CMS Computing Operations

Dirk Hufnagel (FNAL)
for CMS Computing

CHEP, 14.04.2015

Dirk showed how non-CMS sites, and Cloud, Batch, and HPC resources have been made available to CMS opportunistically

Using BOSCO submission mode inside glideInWMS

- Use an ssh tunnel from the factory to submit a pilot to a batch slot at Gordon. All supported out of the box, no custom factory or anything like that. Some BOSCO code needed at Gordon.

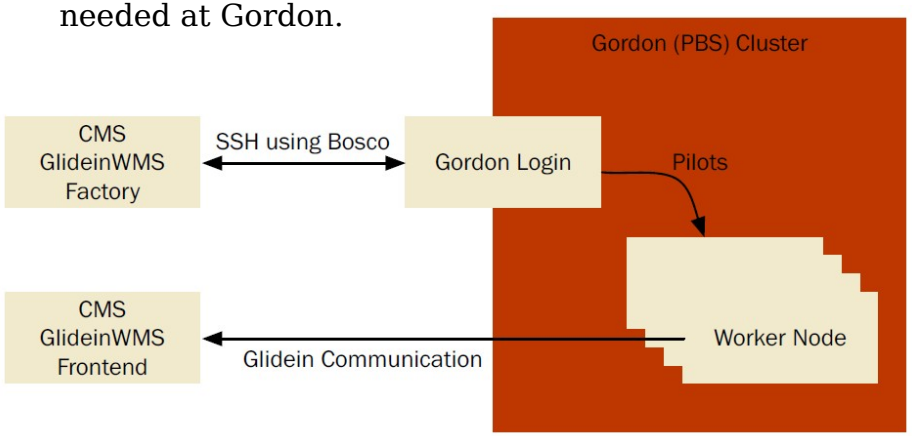

SDSC status

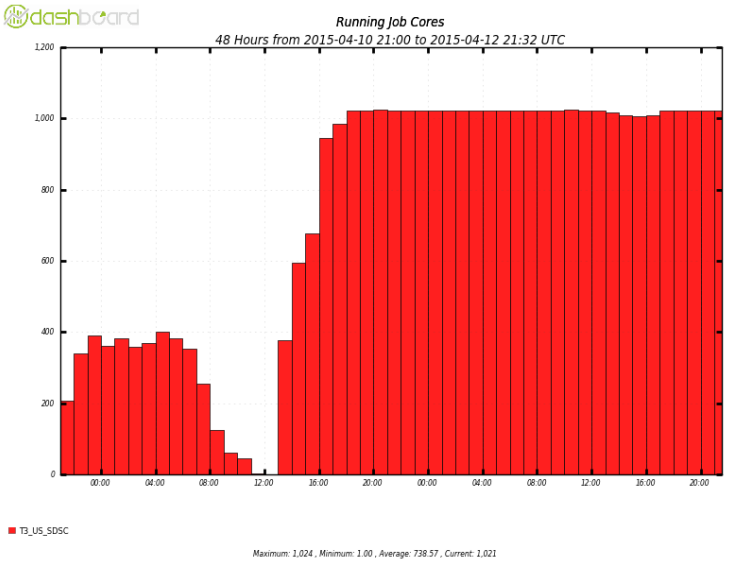

Session 3: Tuesday

Use by experiments

HEP cloud production using the CloudScheduler/HTCondor Architecture

Ian Gable
University of Victoria
on behalf of many people and groups

CHEP 2015 Okinawa, 2015-04-15

Ian explained how CloudScheduler works and how it is being used to provide batch system access to cloud resources for HEP jobs.

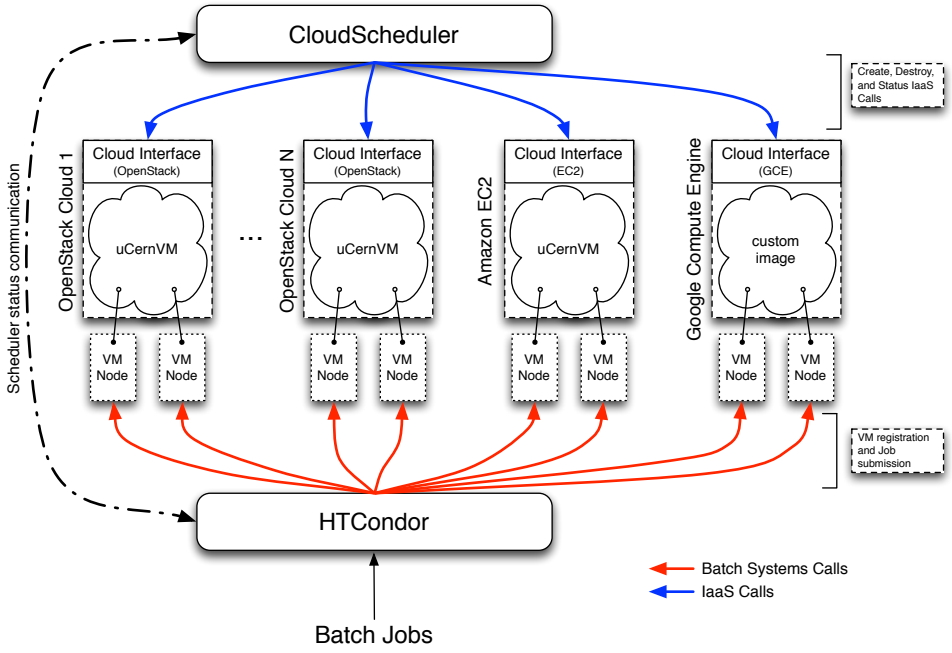

Clouds in use

Currently Operating

Site	HEP Cloud?	ATLAS	BelleII
ComputeCanada-West	no	✔	✔
ComputeCanada-East	no	✔	✔
CERN	yes	✔	
GridPP-Datacentered	yes	✔	
CANARIE-West	no	✔	
CANARIE-East	no	✔	
Amazon EC2	no	✔	
ChameleonCloud-U Texas	no		✔
Google Compute Engine	no		✔

Cloud can come and go with time.

Past clouds include:

NECTAR Australia,
National Research Council Ottawa,
FutureGrid - U Chicago,
Elephant Coud UVic

~ 4000 cores operating today

For comparison, Canadian T2+T1 running ~5000 today

Imperial College London

CMS Diverse use of clouds

David Colling

David went through the various ways CMS is using Clouds & VMs for production, including the CMS HLT farm, via glideinWMS, and Vacuum-based sites.

Imperial College London

How CMS uses Clouds

GlideinWMS used by CMS for job submission to Grids has been modified to also operate with EC2 Clouds. This has been in place for a number of years and works well. However, required modifications to allow for some of the "features" of early OpenStack releases.

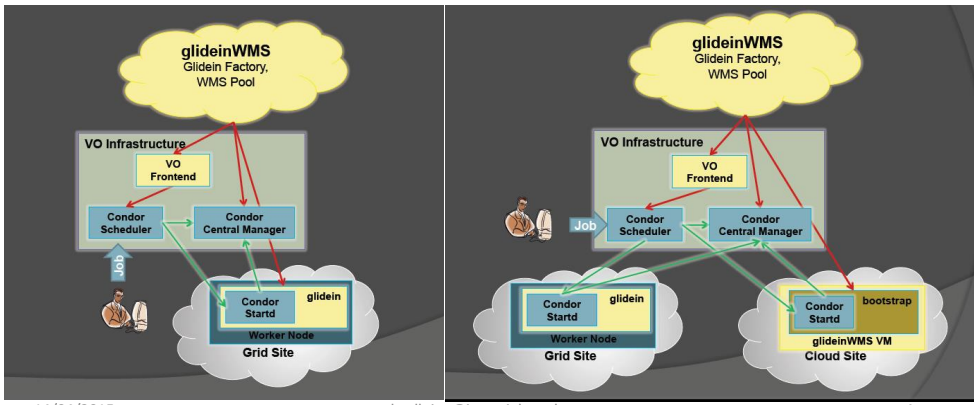

Imperial College London

CMS T0 on AI

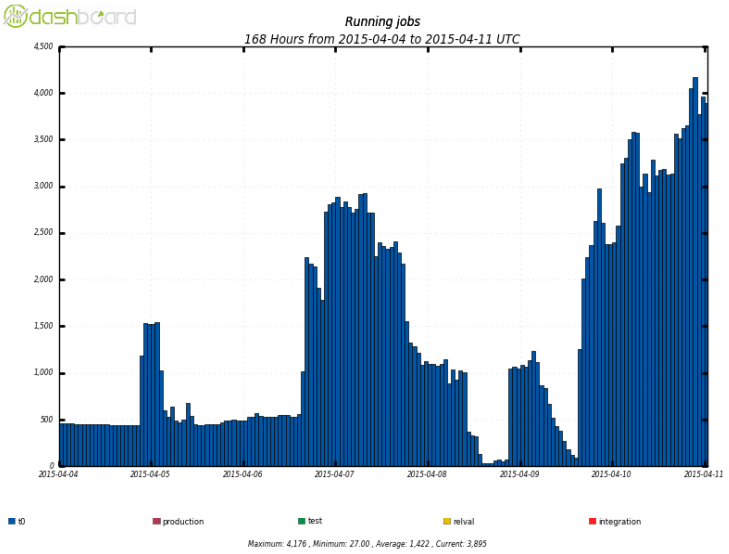

The Evolution of Cloud Computing in ATLAS

Ryan Taylor

University of Victoria

on behalf of the ATLAS collaboration

21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

Ryan went through the different ways in which CloudScheduler (the majority) and the Vacuum platforms are used in ATLAS for production job execution in VMs

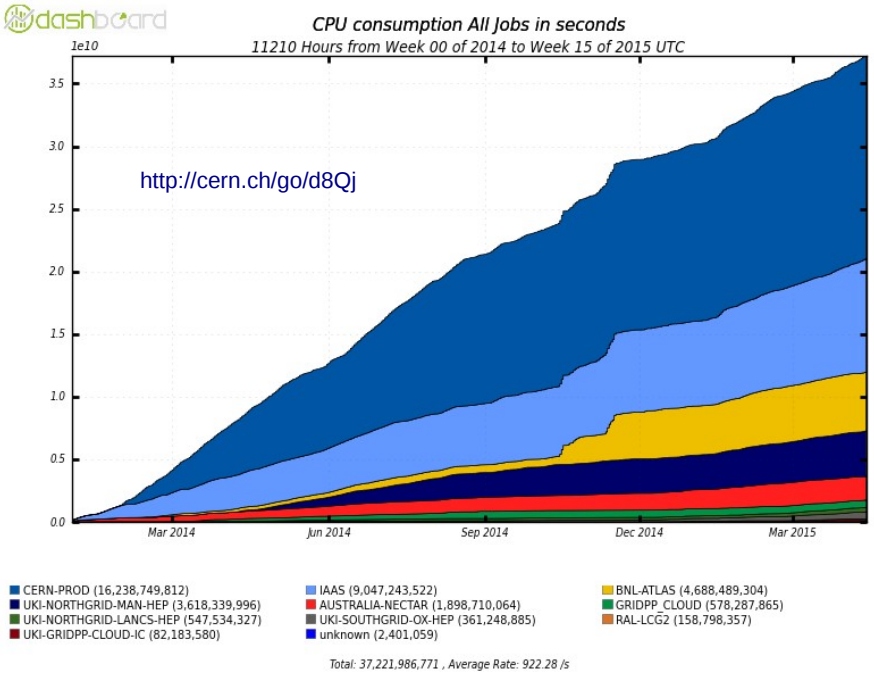

Sim@P1

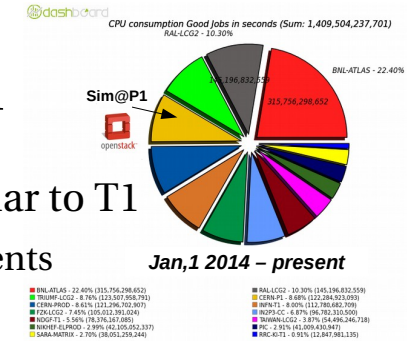

- Resource contribution similar to T1
  - 34M CPU hours, 1.1B MC events
- Used for LHC stops > 24h
- Fast automated switching via web GUI for shifters
  - TDAQ to Sim@P1: 20m (check Nova DB, start VMs)
  - Sim@P1 to TDAQ: 12m (graceful VM shutdown, update DB)
  - Emergency switch to TDAQ: 100s (immediate termination)
- See poster 169

Utilizing cloud computing resources for Belle-II

Randall Sobie
Institute of Particle Physics of Canada
University of Victoria

On behalf of the Belle-II Collaboration

Randall explained how Belle-II's DIRAC based computing system is able to execute jobs on a wide variety of Cloud-based platforms.

Clouds in Belle-II

Clouds at Belle-II member sites

Opportunistic (private and commercial) clouds

Belle-II methods for using clouds

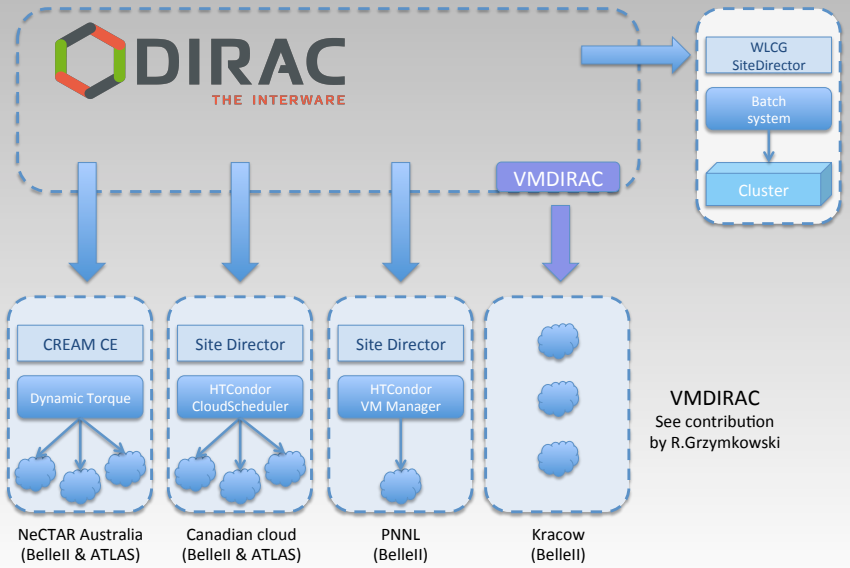

VMDIRAC
See contribution by R.Grzymkowski

LHCb experience running jobs in virtual machines

Andrew McNab,
University of Manchester

Federico Stagni
& Cinzia Luzzi,
CERN

on behalf of the LHCb collaboration

I described the architecture of LHCb's Pilot VMs; and our experience with running production workload at different VM-based sites.

LHCb jobs in VMs

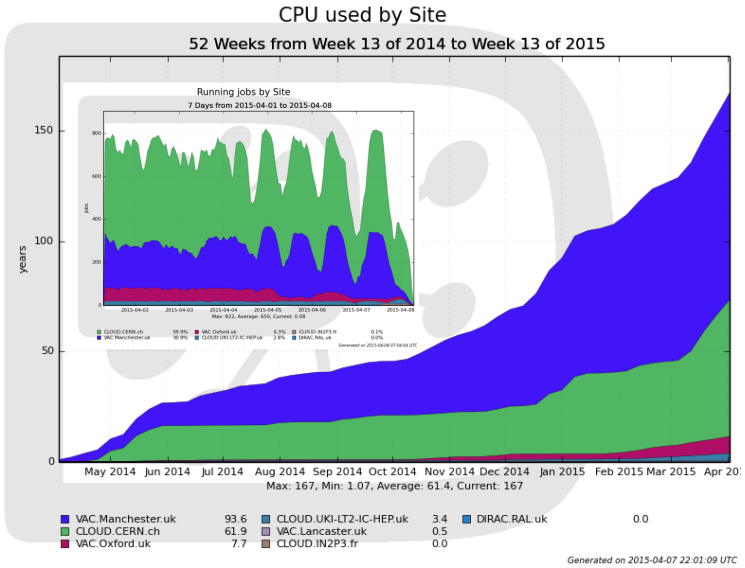

Routine production since May last year.

Increased capacity this year.

Periodic structure due to varying demand from other experiments and in availability of LHCb jobs

Observations

- We get a lot of uniformity from using VMs, and remaining problems tend to be associated with cvmfs use of site's squid cache(s)
- Our use of cvmfs is very sensitive to overloaded squids
  - Shows up as missing/corrupted scripts when starting jobs
  - Need to make sure resource provider is following site config properly
- For a new type of site, ssh access to VMs is very helpful
  - Logging to /etc/machineoutputs also invaluable
  - Need co-operation from resource provider at this stage
- If Vac or Vcycle is already working, tends to be trivial to add LHCb VMs to the site.
- Outcome: improving logging via LHCb-wide Pilot 2.0 to make it easier to debug problems remotely.
- Once the site is working, very little intervention required on either side.

BESIII physical offline data analysis on virtualization platform

Qiulan Huang
huangql@ihep.ac.cn
Computing Center, IHEP,CAS
CHEP 2015

Qiulan's talk outlined the IHEP cloud resources, and explained some detailed work on VM performance optimisation.

IHEPCloud

- Lauched in May 2014.
- A private Iaas platform aiming to provide a self-service cloud platform for users and IHEP scientific computing
- Open for any user who has IHEP email account (>1000 users, >70 active users)

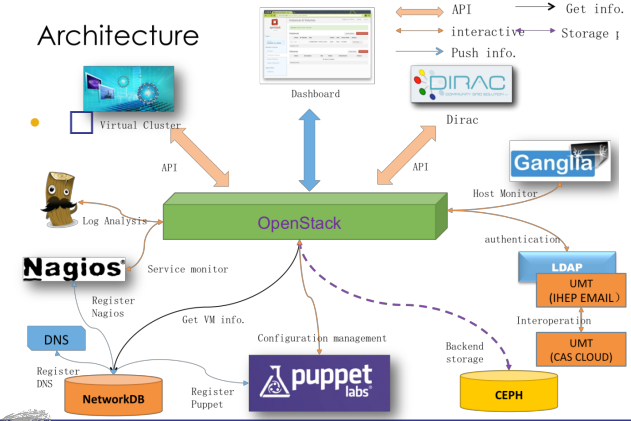

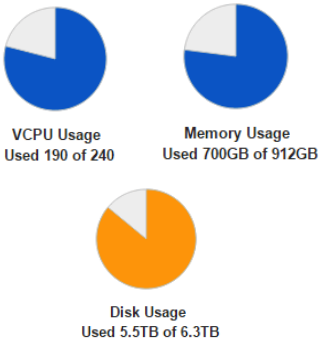

Optimized CPU performance

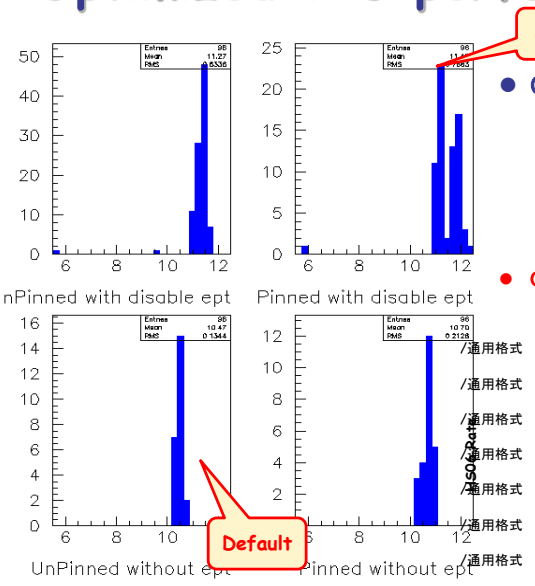

- CPU benchmark testing
  - Specifications
    - Intel(R) Xeon(R) CPU X5650 (2.67GHz) , 8 CPU cores , 24GB内存
    - OS : SLC release 5.5 (Boron) 2.6.18-194.11.3.el5.cve20103081 , 64bit
    - KVM-83
  - Tools
    - HEP-SPEC06
- Optimized CPU performance increased about 3%

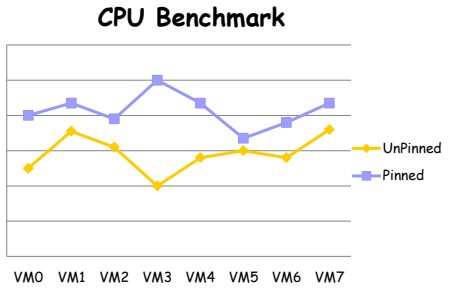

Yandex School of Data Analysis

SHiP
Search for Hidden Particles

Skygrid

Alexander explained how SkyGrid schedules work across several platforms, including tasks translated into Javascript for execution in web browsers.

Alexander Baranov, Konstantin Nikitin, Andrey Ustyuzhanin

Skygrid architecture

«Job → Queue → Worker» scheme:

- Job = JSON description of what to do
- Queue = Metascheduler service (matches jobs to resources)
- Worker = Basically anything (single computer, cluster, browser, ...)

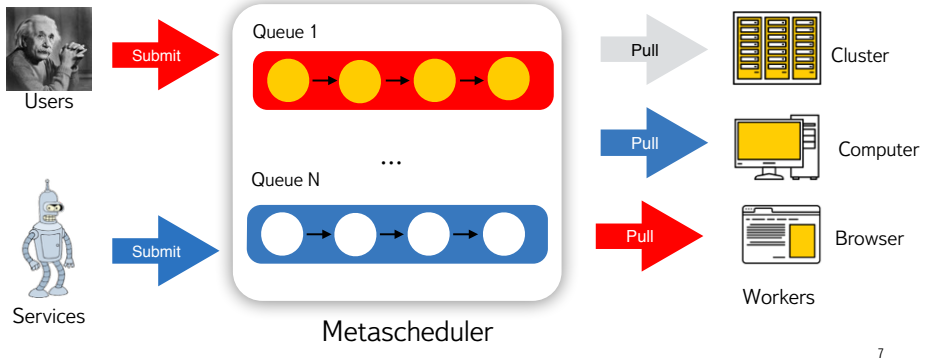

In-browser volunteer computing

- Skygrid provides API so your browser can pull a job
- Helpful for volunteer computing
- Prototype: generate MC in browser (Pythia→ LLVM → JavaScript)
- Details at poster session on DiBroCop — Distributed inBrowser ComPutations

Seriously, you can run it on your phone:

Session 4: Thursday

Technologies

Status and Roadmap of CernVM

G Ganis

For the CernVM team

16 April 2015
CHEP 2015, Okinawa, Japan

Gerardo explained how CernVM is continuing to evolve to support its wide range of uses in HEP, including with the addition of containers

CernVM Users 2015

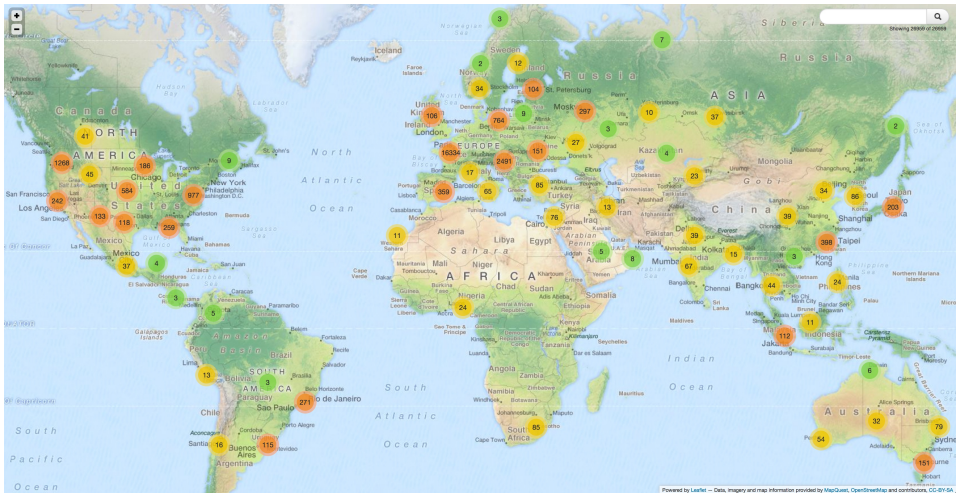

≈ 27000 distinct IP addresses

CernVM as container

Roadmap
1: Make CernVM-FS repositories available inside a container
2: Writable overlay via union file system

Options for CernVM-FS (1)

1: Fuse, mapped from host
- Shared cache
- Promising prototype in progress (➔D Berzano's talk)

2: Using Parrot CernVM-FS
- Pure user-space (ptrace)
- User-mode CernVM demonstrated (➔next slide)
- Ongoing work on performance and stability

Root file system (/) layout

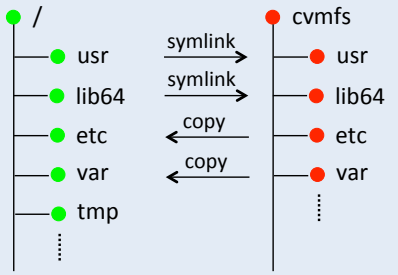

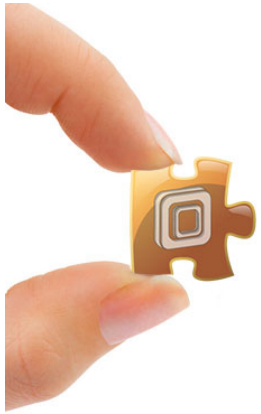

CernVM WebAPI

Controlling VMs from the web

I.Charalampidis, J.Blomer, D.Berzano,
P.Buncic, G.Ganis, R.Meusel (CERN)

CHEP, Okinawa, 16 April 2015

Ioannis showed how the CernVM WebAPI makes it easy to install VMs from a web browser

CernVM WebAPI

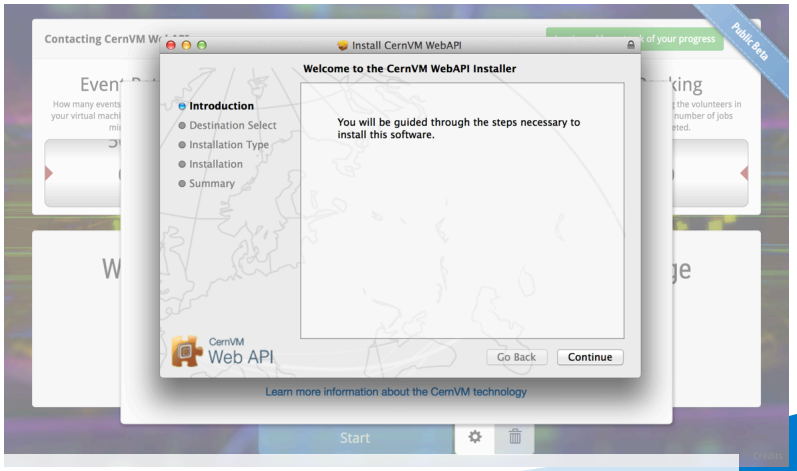

Installation in a pop-up window

The CERN60 Challenge

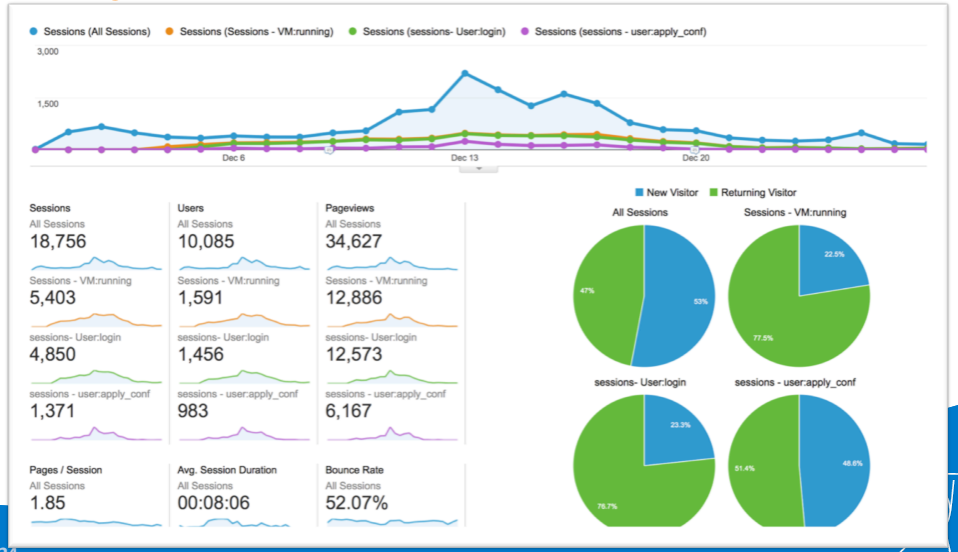

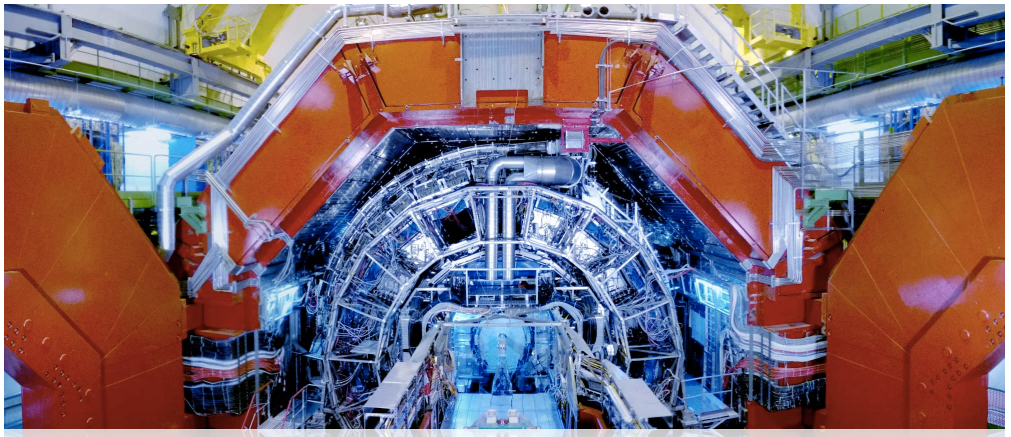

Lightweight scheduling of elastic analysis containers in a competitive cloud environment: a Docked Analysis Facility for ALICE

Dario Berzano
ALICE Offline - CERN

Computing in High Energy and Nuclear Physics - Apr 13-17, 2015

Dario presented the ALICE model for Pilot Containers, and proposed an extension to Vac to support them

The pilot container

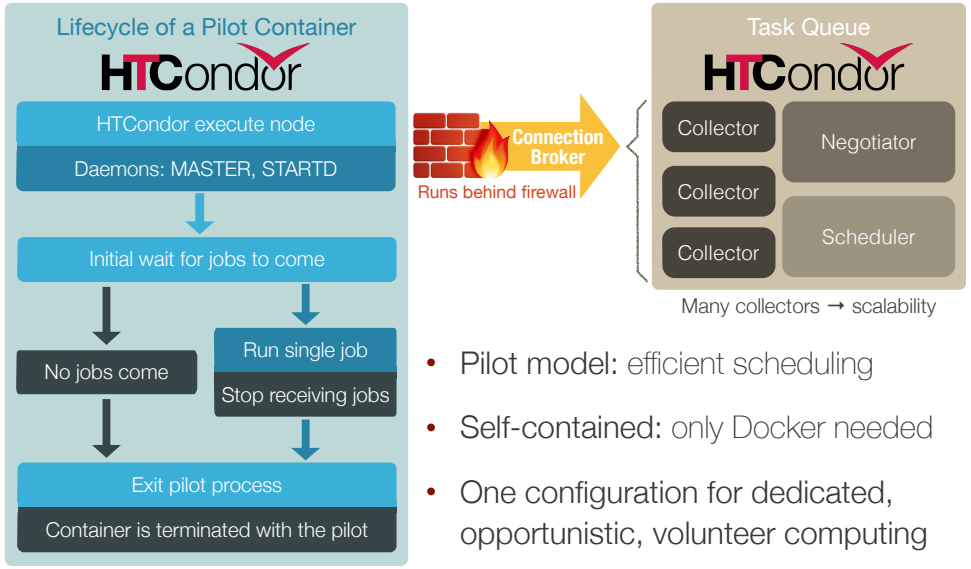

- Pilot model: efficient scheduling
- Self-contained: only Docker needed
- One configuration for dedicated, opportunistic, volunteer computing
- Prototype and doc: github.com/dberzano/cernvm-alice-docker

The container factory

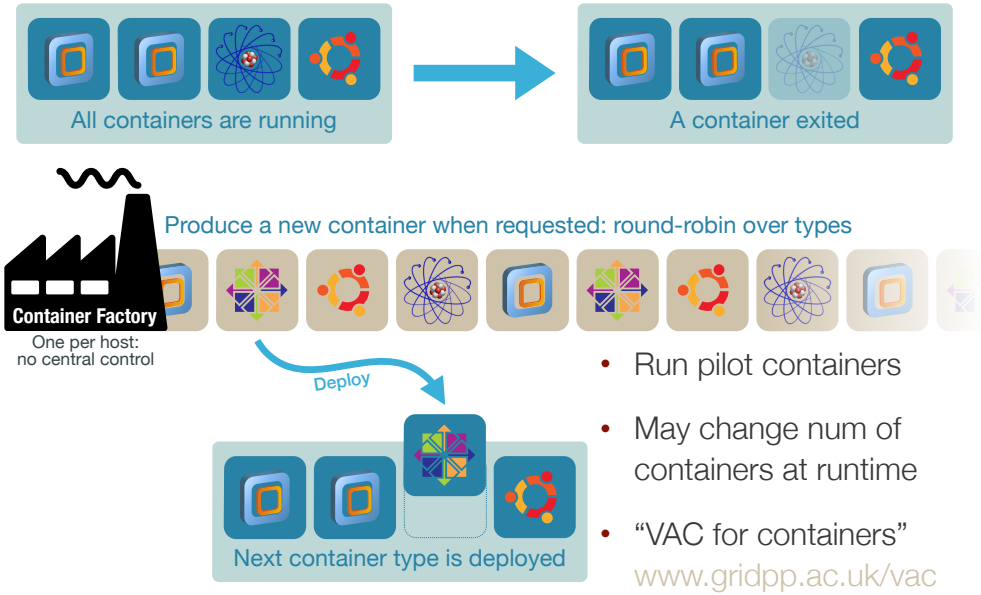

- Run pilot containers
- May change num of containers at runtime
- "VAC for containers" www.gridpp.ac.uk/vac

Cloud federation

Prelude to Hybrid Clouds

- CHEP 2015
- Okinawa, Japan

Marek Denis
CERN
Geneva, Switzerland

CERN openlab

Background image: Shutterstock

Marek explained how identity federation is now possible between OpenStack sites

OS-FEDERATION timeline

OpenStack Summit (November 2013)

Basic concept and initial discussions during design sessions

OpenStack Icehouse (April 2014)

Server-side OS-FEDERATION delivered (located in the extensions namespace)

OpenStack Juno (October 2014)

OS-FEDERATION marked as stable. Client code integrated with official OpenStack libraries and CLIs (CERN uses OS-FEDERATION internally since September 2014)

OpenStack Kilo (April 2015)

Added WebSSO support in keystone. Mapping engine enhancements

Federated authN & authZ

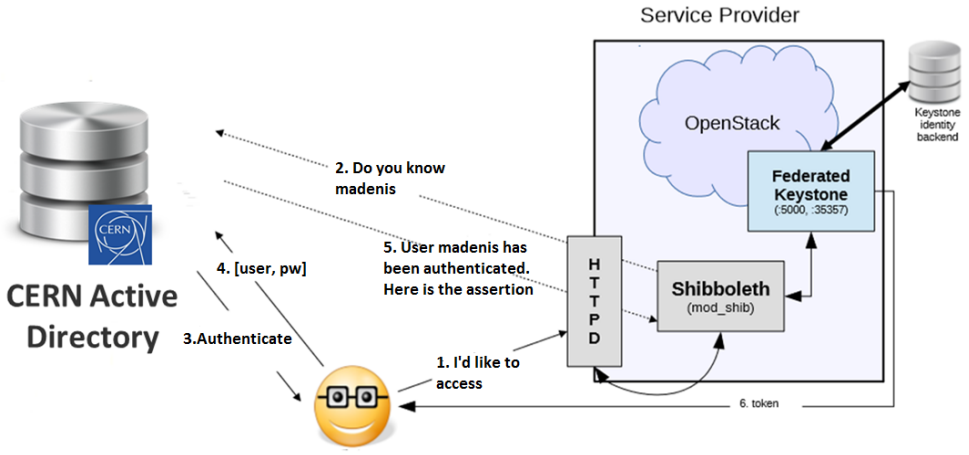

Credits Luca Tartarini

Background image: Shutterstock

Stefano Bagnasco, Domenico Elia, Grazia Luparello, **Stefano Piano**, Sara Vallero, Massimo Venaruzzo
For the **STOA-LHC** Project

INTEROPERATING CLOUD BASED VIRTUAL FARMS

Stefano told us how cloud resources at INFN sites have been federated, both for job execution and data access

21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

THE INFRASTRUCTURE

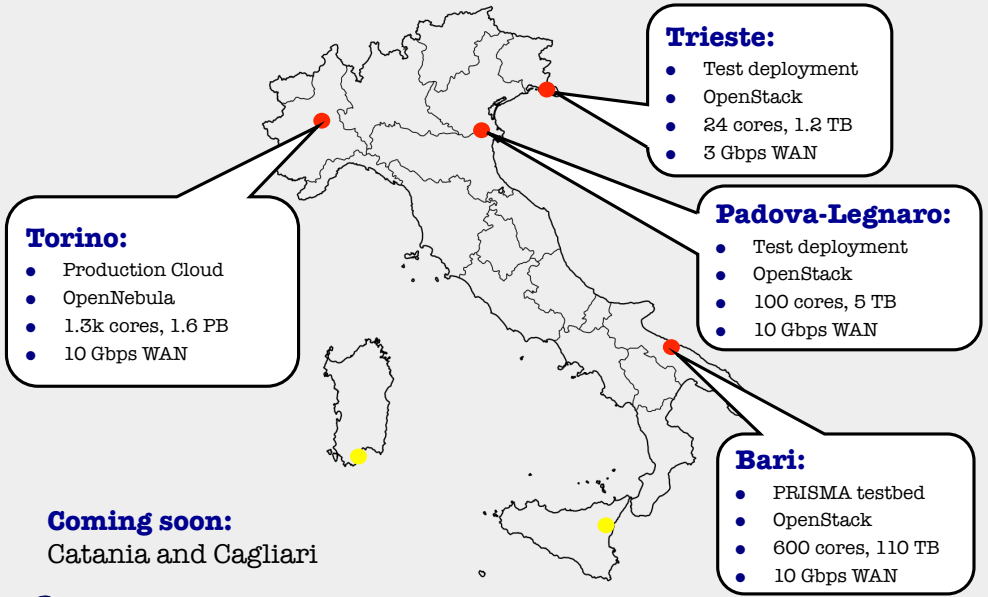

KEY COMPONENT: THE VAF

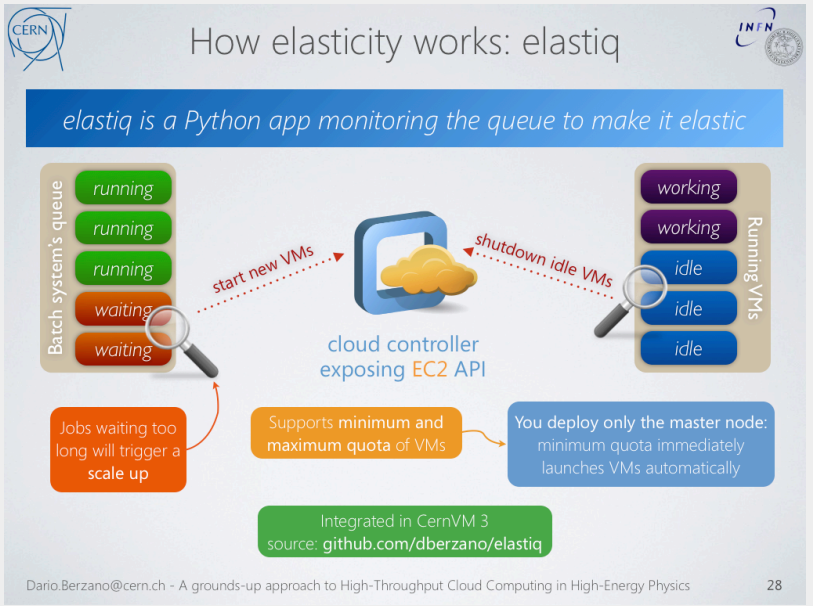

Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Cloud Services for the Fermilab Scientific Stakeholders

Steven Timm[1], Gabriele Garzoglio[1], Parag Mhashilkar[1], Joe Boyd[1], Gerard Bernabeu[1], Neha Sharma[1], Nicholas Peregonow[1], Seo-young Noh[2], Hyun Woo Kim[1], Sandeep Palur[3], Ioan Raicu[3]

Fermilab[1], KISTI[2], Illinois Institute of Technology[3]

Parag described how Fermilab has demonstrated extending their cloud capacity in response to demand, using Amazon resources

Running AWS Jobs as function of time, Oct 23. 2014

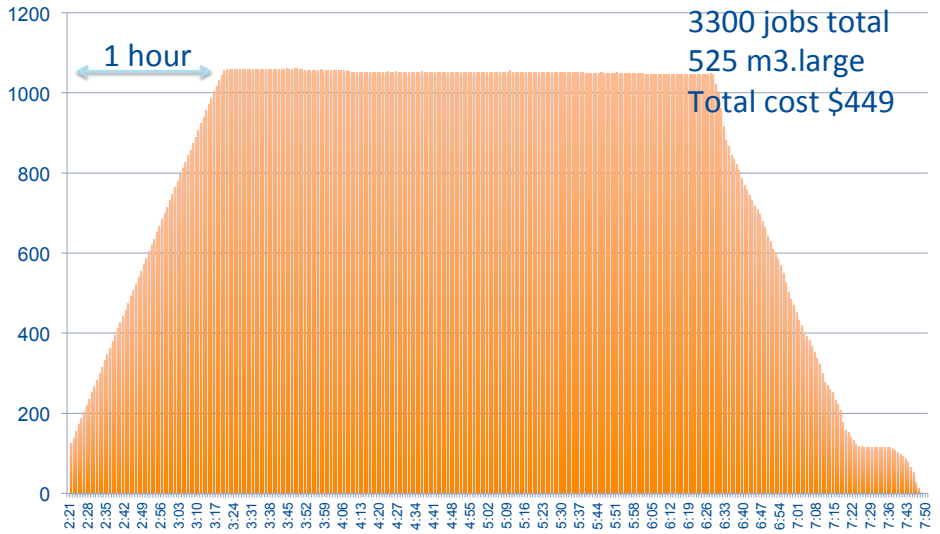

Hybrid Cloud: FermiCloud & Amazon Web Service

- Use FermiCloud first & Pay AWS only during spikes
  - Need to make differences in cloud providers transparent
    - User jobs don't notice any difference
    - Solution should be manageable by the operations

Have prototyped using Amazon Route 53 service to have a single service IP

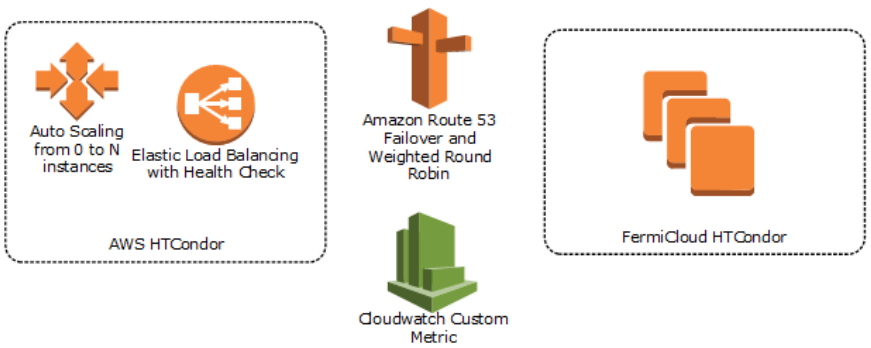

Summary

- 27 talks presented in 4 sessions
- This talk didn't have time to cover the 31 posters in depth
  - It is well worth looking through them in Indico
- The sessions were very well attended:
  - People were standing/sitting on the floor at some point during each session in a room with 70+ chairs
- The sessions were lively and interesting, and a lot of new ideas exchanged
  - Discussion component was vital
- We're looking forward to reading the papers next!

Gartner Hype Cycle