Clouds in High Energy Physics

Randall Sobie

University of Victoria

Overview

- Clouds are integral part of our HEP computing infrastructure
  - Primarily Infrastructure-as-a-Service (IAAS)
- Our use is wide ranging and diverse
  - CERN Agile Infrastructure
  - Tier-1 computing at centres such as BNL, FNAL and RAL
  - Tier-2 computing around the world
- Expanding use of HEP-clouds, private clouds and commercial clouds

Motivation

A wide range of reasons for using clouds

- Ease management of existing infrastructure
- Separation of application and system administration
- Simplifies allocation of resources
- Leverage software development
- Opportunistic computing
  - Non-HEP computing centres
  - Commercial cloud resources

Types of cloud resources

Cloud computing in HEP is typically providing 5-20% of the processing of current projects

"Dedicated" clouds
(Owned by HEP)

"Opportunistic" clouds
(private and commercial)

Cloud deployments

Traditional
bare-metal

Static cloud
(e.g.. LTDA BaBar, HLT clouds)

Standalone/private
cloud
(e.g. PNNL, NorduGrid)

Distributed clouds
(e.g. UK, Canada,
Australia, INFN Clouds)

Bare-metal or in-house cloud with external cloud
(e.g.. CERN, BNL)

OpenStack Clouds at CERN

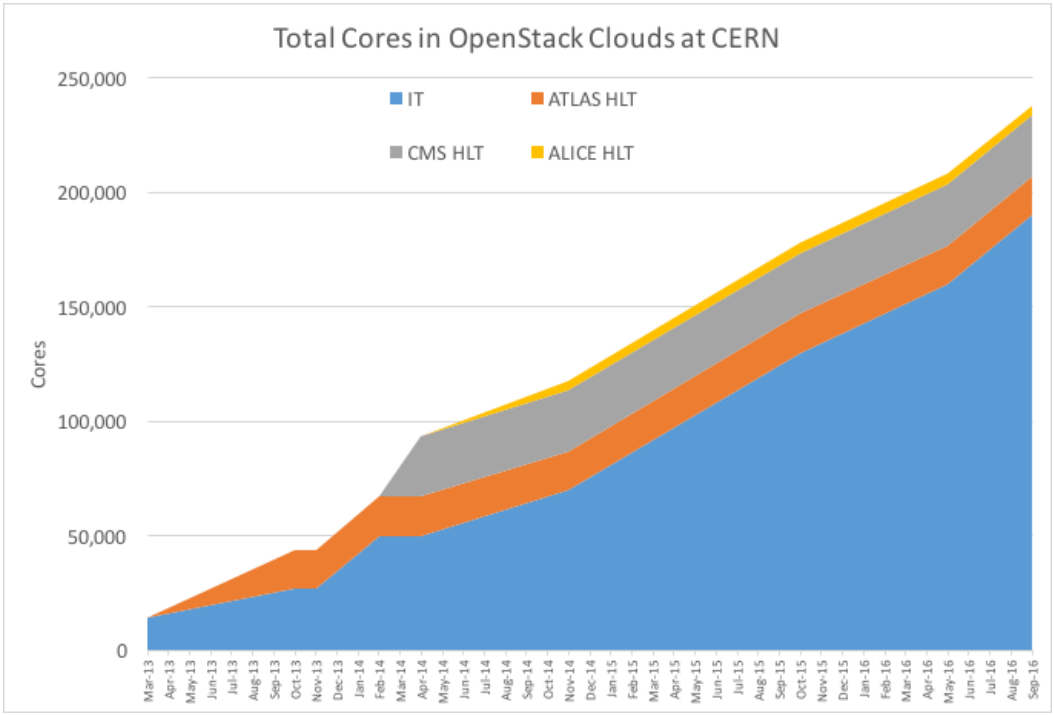

In production:

- 4 clouds
- >230K cores
- >8,000 hypervisors

>90% of CERN compute resources are virtualized

Up to 42K cores to be installed in the next few months subject to funding

Clouds at RAL

Science & Technology Facilities Council

OpenNebula

- 892 cores utilizing a Ceph storage backend.
- 3 alternating racks of CPU and Storage nodes.
- Tier 1 services now running on Cloud VMs.
- Engaging with various European Cloud projects (e.g DataCloud).

S3 and Swift Storage

- Storing Docker images for Container Orchestration via Swift.

Openstack service under development. Available to LHC VOs next year.

openstack CLOUD SOFTWARE

Batch Work on the Cloud

- For ~1.5 years the RAL HTCondor batch system has made opportunistic use of unused cloud resources.
- HTCondor rooster daemon used to provision VMs.
- Running jobs from all 4 LHC experiments & many non-LHC experiments.

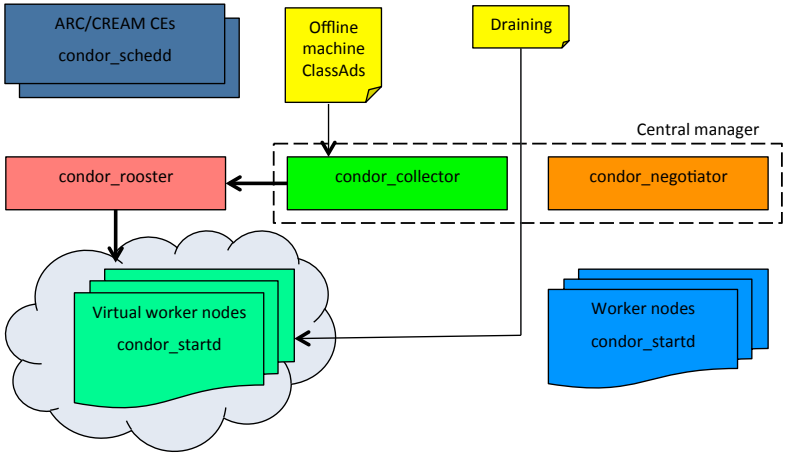

FNAL/CMS AWS January/February 2016

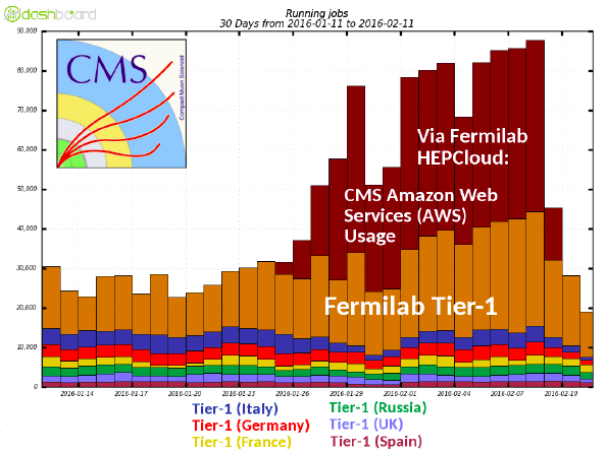

13

Tier-1
Cloud bursting onto EC2

BNL/ATLAS AWS September 2015

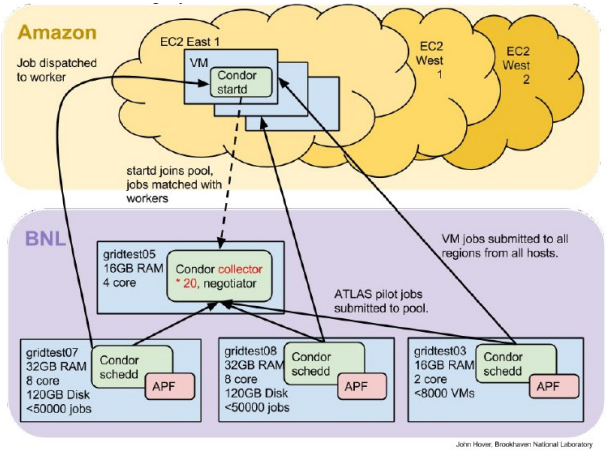

9

Special purpose clouds
BaBar Long Term Data Access (LTDA) System

Ability to preserve data and analysis capability for BaBar
(stopped data taking in 2008)

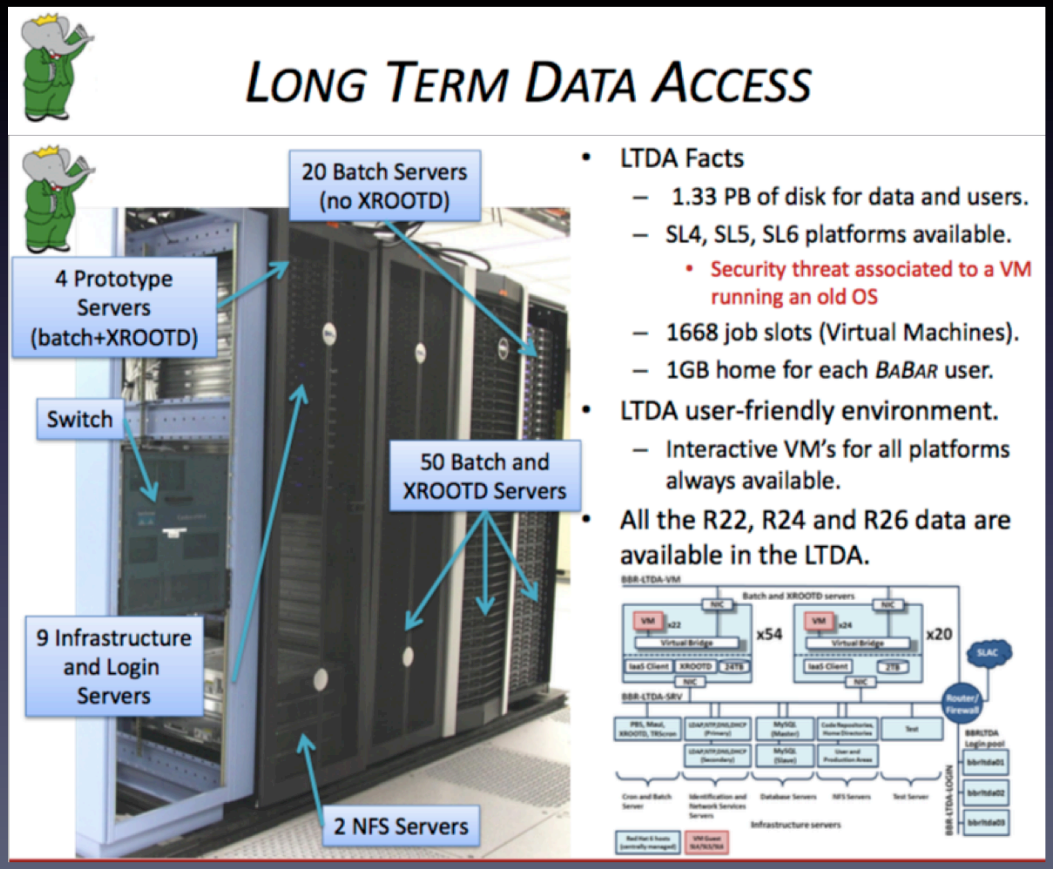

High level trigger farms of the LHC Experiments

(large multi-10K core systems)

Virtual machines are booted during no-beam periods

Examples of Tier-2 cloud deployments

Australian Belle II Grid Site

Single CREAM CE services
ATLAS Tier-2 (Torque)
and
Belle II site (Dynamic Torque)

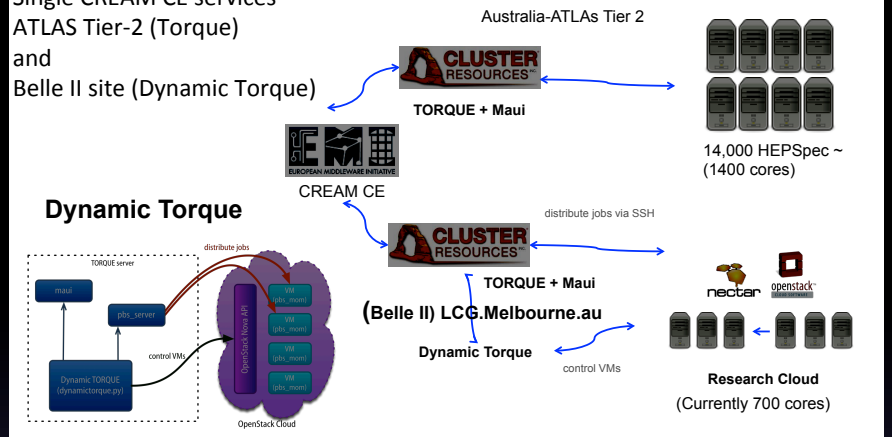

SWITCHengines – Swiss NREN commercial cloud

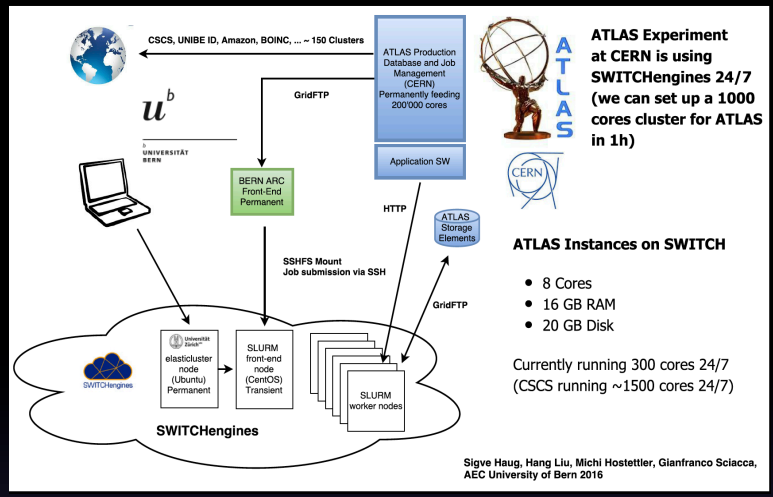

NorduGrid Secure hybrid cloud

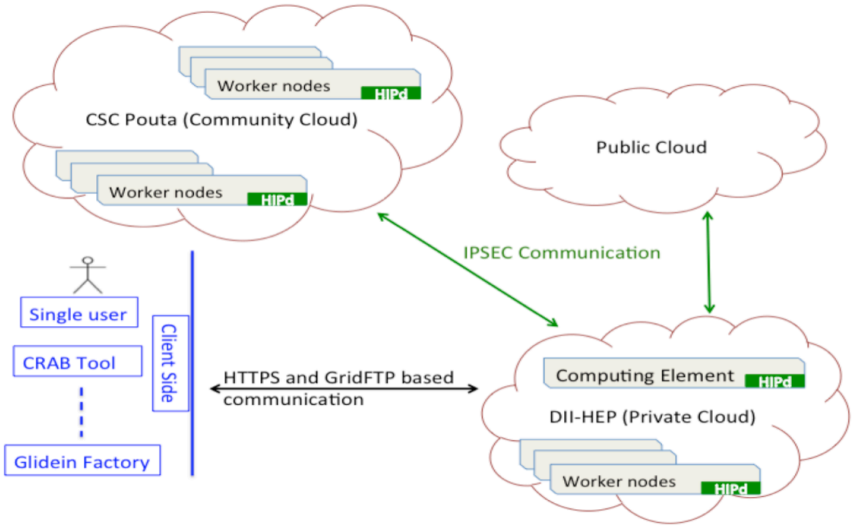

Why private cloud?

Pacific Northwest NATIONAL LABORATORY
Proudly Operated by Battelle Since 1965

- Chosen for flexibility, efficient use of compute resources for services
- Provides easy load-balancing and availability features
- Provides templating features
- Easy re-use of templates to test and instantiate new server instances
- Non-systems staff can provision their own instances of services
- Software Defined Networking is more malleable than physical networking, encourages better networking practices, including security

UK / GridPP

- Clouds at HEP institutions (Oxford/Imperial).
- ECDF cloud in Edinburgh has recently made available to the HEP
- UK Vacuum deployment
- Commercial cloud – DataCentred Openstack cloud

Italy / INFN

- PrivateOpenStack Cloud (Padova-Legnaro) called CLOUD AREA PADOVANA
- ~ 25 user groups/project
- CMS production

PNNL / Washington

Private OpenStack cloud for Belle II project (KEK) and other local users

Canada

Distributed cloud system for ATLAS and Belle II

10-15 clouds
HTCondor/CloudScheduler
4000-5000 cores

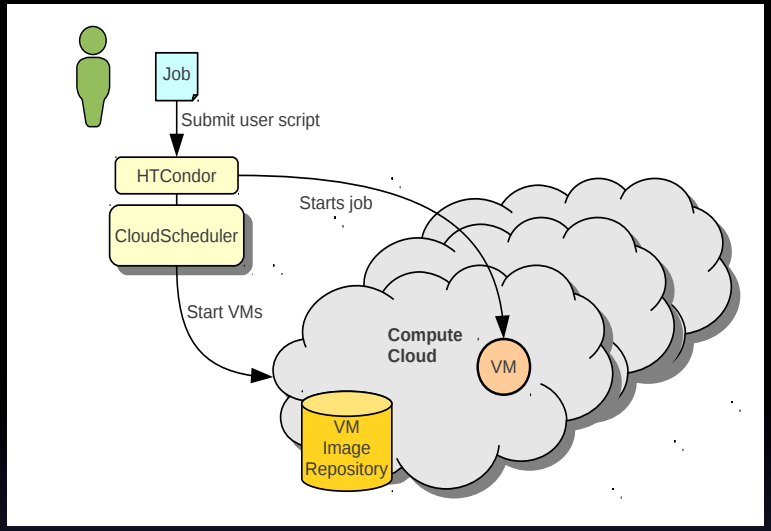

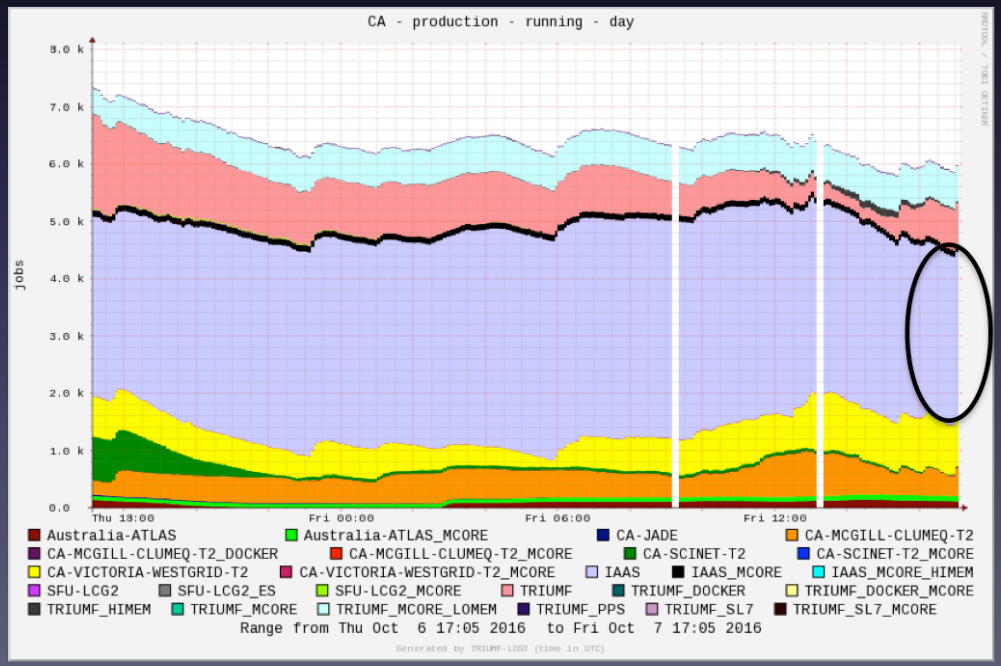

ATLAS jobs on cloud for CA-system
10 clouds
4300 cores

Job scheduling/VM provisioning

- Variety of methods for running HEP workloads on clouds
  - VM-DIRAC (LHCb and Belle II)
  - VAC/Vcycle (UK)
  - HTCondor/CloudScheduler (Canada)
  - HTC/GlideinWMS (FNAL), HTC/VM (PNNL), HTC/APF (BNL)
  - Dynamic-Torque (Australia)
  - Cloud Area Padovana (INFN)
  - ARC (NorduGrid)
- Each method has its own merits and often was designed to integrated clouds into an existing infrastructure (e.g. local, WLCG and experiment)

Commercial clouds

- Amazon EC2 and Microsoft Azure
  - Short-term multi-10K tests
  - Long-term 1K-scale production
- GCE evaluation but no production
- Other commercial OpenStack clouds
  - DataCentred (UK), SWITCHengines (Switzerland)
- CERN commercial cloud procurement

Network connectivity

- Amazon and Microsoft clouds are connected to the research networks in North America (probably GCE as well)

- Trans-border or trans-ocean traffic can be an issue
  - Become an important discussion topic in the LHCONE meetings

- Private opportunistic clouds
  - traffic flows over research network but not LHCONE network

CPU Benchmarks

New suite of "fast" benchmarks

- HEPiX Benchmark Working Group
- Suite available includes "fast HS" (LHCb) and Whetstone benchmarks
  - Write to ElasticSearch DB
- Run benchmarks in the pilot job or during the boot of the VM

Data storage

- Data written to local storage on node and then transferred to selected SE
- UK group has done some work integrating their object store with ATLAS
- BNL using S_3 storage on EC_2 for T2-SE

Monitoring

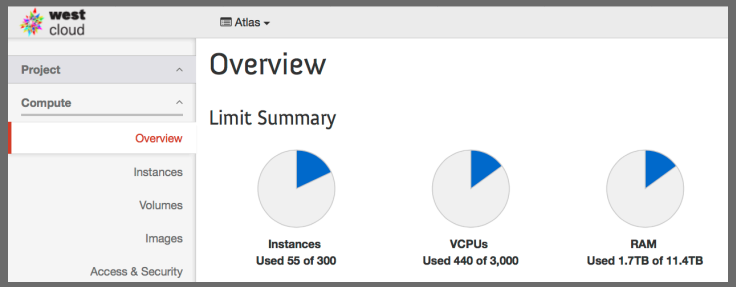

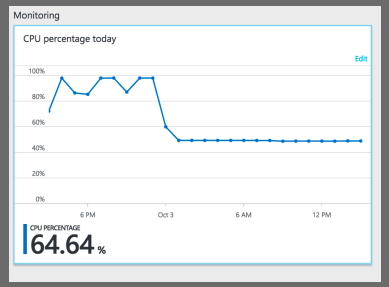

Cloud or site monitor

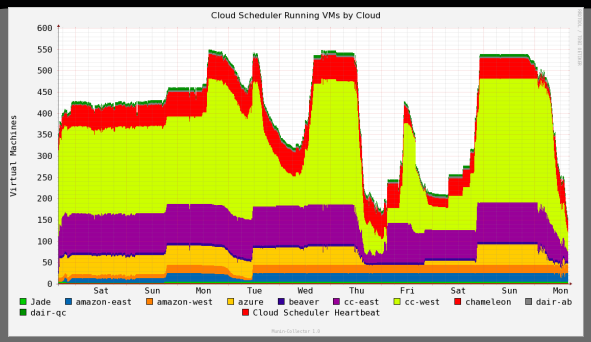

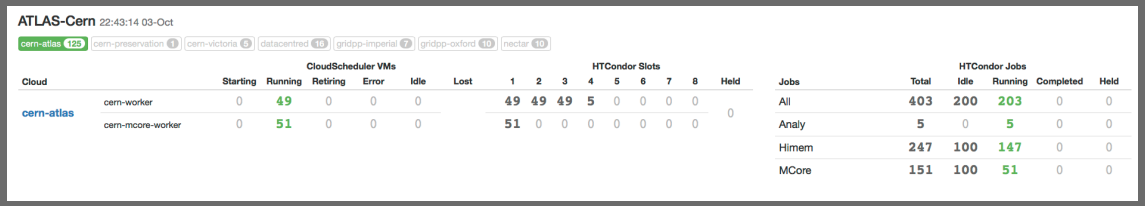

Cloud System monitor

Sensu, Munin, RabbitMQ, Mongo-DB

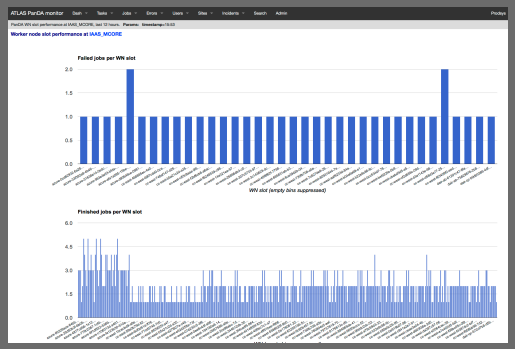

Application monitor

Panda

MONTH

Cloud	#	Bmk	User	Total
				Total
beaver	18	15.6	49.5	60.5
cc-west	1610	19.0	3228.4	4205.4
cc-east	296	14.5	931.1	1129.1
chameleon	136	21.1	976.1	1269.5
dair-ab	2	12.7	30.6	32.9
dair-qc	7	12.3	85.3	92.2
azure	126	21.4	1518.2	1800.3
ec2	106	9.8	36.4	184.3
Total			6855.6	8774.2

Monday October 03 15:00:01

Benchmarks and accounting

ElasticSearch DB

Summary

- Clouds in HEP
  - Growing, diverse use of clouds
  - Typically integrated into an existing infrastructure
  - Seen as a way to better manage multi-user resources
- Opportunistic research clouds
  - Easy way to utilize clouds at non-HEP research computing facilities
  - No requirement for on-site application specialists or complex software
- Commercial clouds
  - EC2/Azure as well as other OpenStack clouds
  - Trans-border network connectivity challenges