





Cloud Status

Laurence Field
IT/SDC

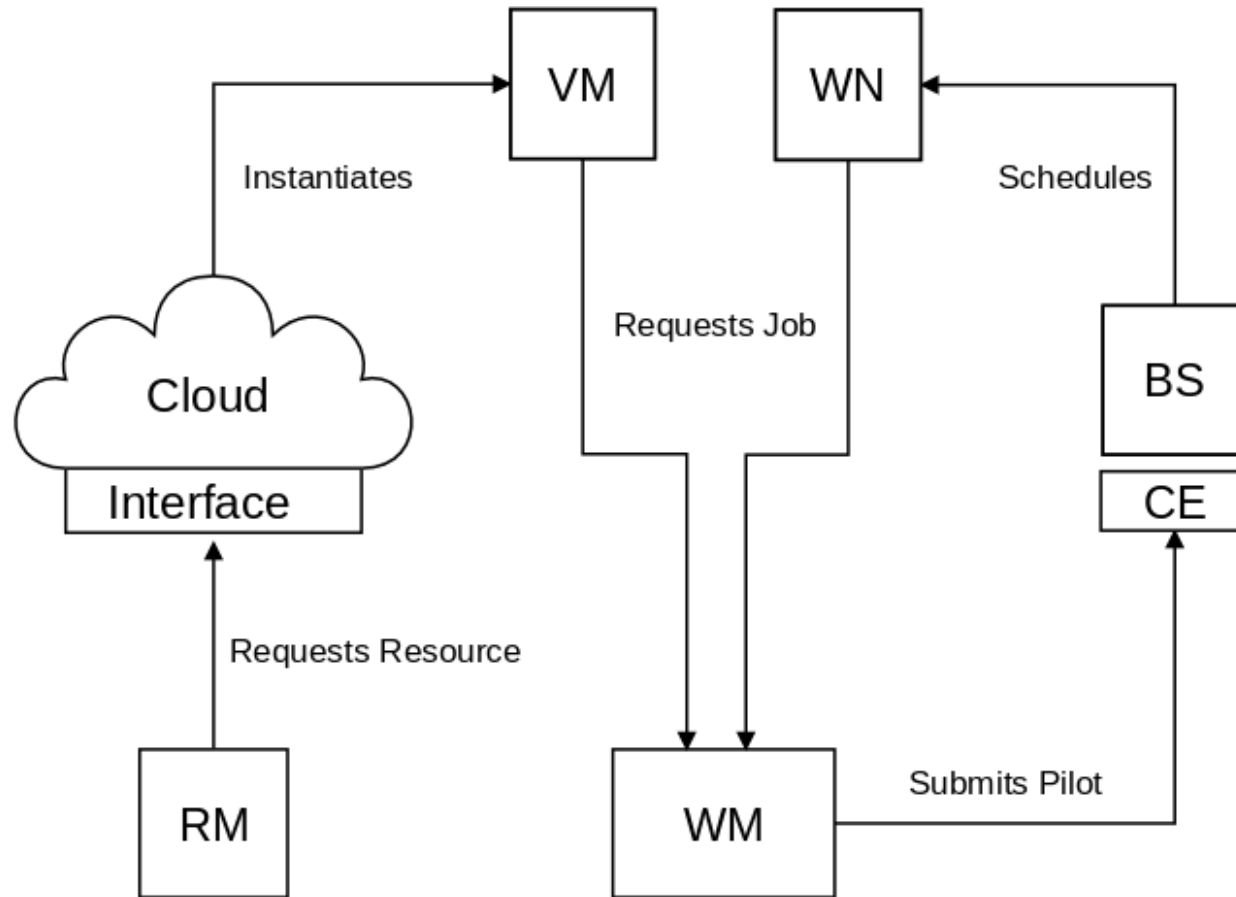
09/09/2014



Cloud



High Level View



Areas

- Image Management
- Capacity Management
- Monitoring
- Accounting
- Pilot Job Framework
- Supporting Services

Image Management

- Provides the job environment
 - Software
 - CVMFS
 - PilotJob
 - Configuration
 - Contextualization
- Balance pre- and post-instantiation operations
 - Simplicity, Complexity, Data Transfer, Frequency of Updates
- Transient
 - No updates of running machines
 - Destroy (gracefully) and create new instance

CernVM

- The OS via CVMFS
 - Replica via HTTP a reference file system
 - Stratum 0
- Why?
 - Because you needed CVMFS
 - Removes the overhead of distributed image management
 - Manage version control centrally
- CernVM as a common requirement
 - Availability becomes an infrastructure issue
 - Recipe is the responsibility of the VO
- The goal is to start a CernVM-based instance
 - And provide the minimal contextualization

Capacity Management

- Because managing a VM isn't the focus
 - It is about ensuring there is enough resources (capacity)
- Requires a specific component with some intelligence
 - Do I need to start a VM and if so where?
 - Do I need to stop a VM and if so where?
 - Are the VMs that I started OK?
- Existing solutions focus on deploying applications in the cloud
 - Different components, one cloud
 - May manage load balancing and failover
 - Is this a load balancing problem?
 - One configuration, many places, enough instances?
- Developing our own solutions
 - Site centric
 - The VAC model
 - VO centric

Tooling

- VAC
 - Site-controlled bare metal provisioning
- V-Cycle
 - VAC for OpenStack tenants
- Cloud Scheduler
 - Condor-based VM control
- Glidein WMS
 - Adaptation of the CondorHTC pilot job approach
- BOINC
 - Extra capacity for certain classes of job

Accounting



- Helix Nebula
 - Pathfinder project
 - Development and exploitation
 - Cloud Computing Infrastructure
 - Divided into supply and demand
 - Three flagship applications
 - CERN (ATLAS simulation)
 - EMBL
 - ESA
- **FW: New Invoice!**
 - *Can you please confirm that these are legit?*
 - Need to method to **record** usage to cross-check invoices



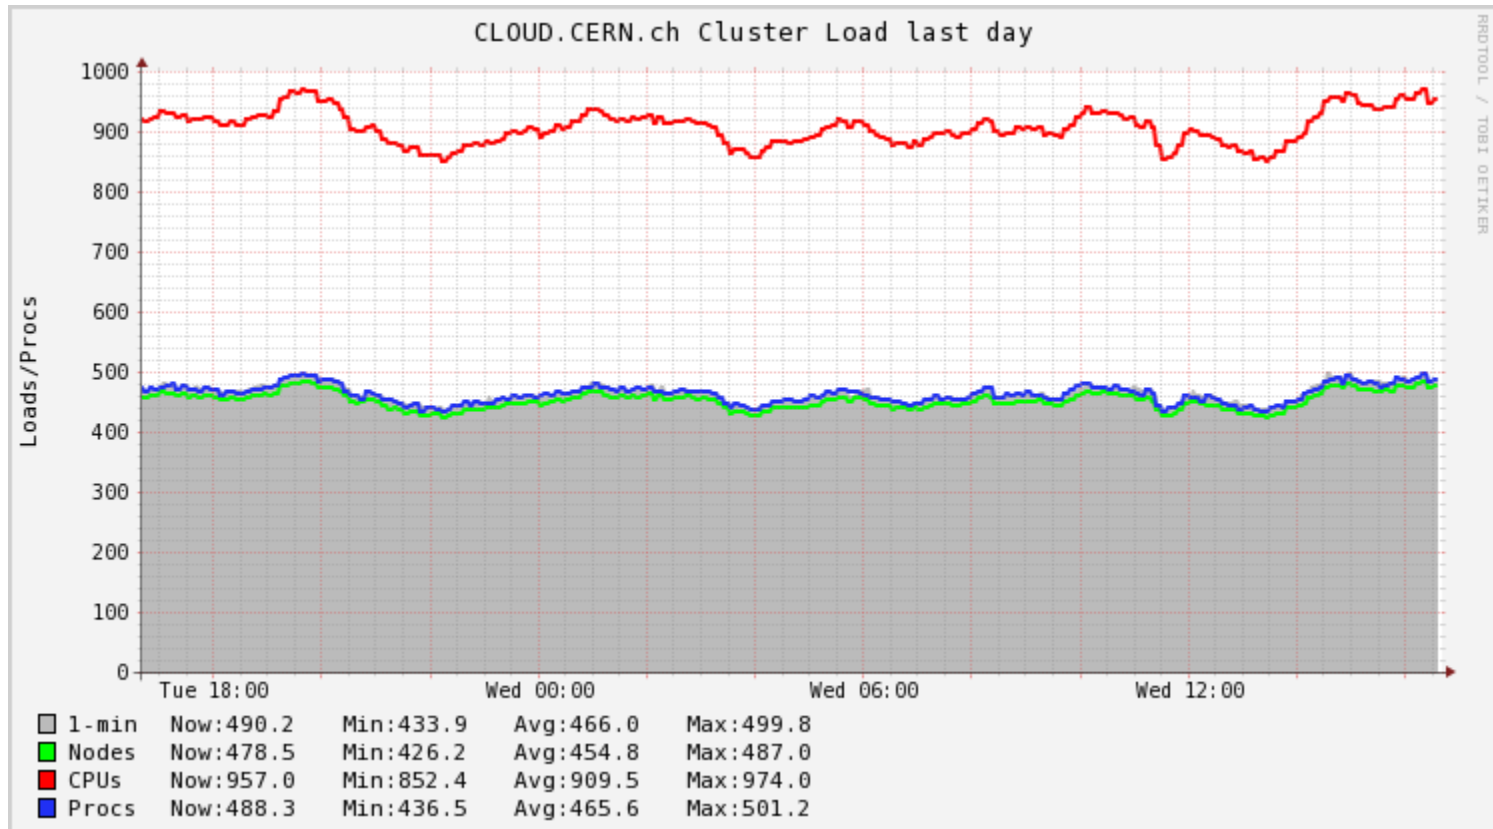
Consumer-Side Accounting

- Monitor resource usage
 - Course granularity acceptable
 - No need to accurately measure
- What, where, when for resources
 - Basic infrastructure level
 - VM instances and whatever else is billed for
- Report generation
 - Mirror invoices
 - Use same metrics as charged for
- Needs a uniform approach
 - Should work for all VOs
 - Deliver same information to the budget holder

Ganglia-based Accounting

- Ganglia already adopted to monitor clouds
 - Simple, scalable, customizable, just works
 - Used by ATLAS, LHCb, CMS
 - ALICE use MonaLisa instead
- Default Metrics
 - Total CPUs (cores)
 - Hosts (VM instances)
 - Load (CPU efficiency)
 - Network input/output
 - Local disk size
- Resource Accounting Record
 - Take the integral

Standard Ganglia Plot



Accounting Records

WLCG Cloud, Usage

Cloud Resources ▾

VO: MONTH: YEAR:

Note: Sorting by multiple columns at the same time can be activated by 'shift' clicking on the column headers which they want to add to the sort. Hovering mouse over the column headers to get descriptions of table columns.

Site	Wall Duration(h)	CPU Duration(h)	CPU Count	Network Inbound(Gb)	Network outbound(Gb)	Memory(GB)	Disk(GB)	Cloud Type
CLOUD.CERN.ch	30352.70	16355.48	420.62	4654.52	393.60	1410.00	0	
VAC.Manchester.uk	5665.28	5761.42	78.59	1399.09	105.73	336.34	0	
CLOUD.UKI-LT2-IC-HEP.uk	147.72	153.93	6.13	68.33	5.45	6.96	0	
VAC.Oxford.uk	92.25	97.98	1.92	25.21	1.61	4.50	0	
VAC.Lancaster.uk	71.39	79.87	2.96	22.95	1.31	3.43	0	
VAC.UKI-LT2-IC-HEP.uk	0	0	0	0	0	0	0	
CLOUD.PIC.es	0	0	0	0	0	0	0	
CLOUD.CERNMP.ch	0	0	0	0	0	0	0	
CLOUD.CERN8.ch	0	0	0	0	0	0	0	
BOINC.World.org	0	0	0	0	0	0	0	
BOINC.Manchester.uk	0	0	0	0	0	0	0	
BOINC.CERN.ch	0	0	0	0	0	0	0	
Total:	36329.34	22448.68	510.22	6170.10	507.70	1761.23	0.00	

Showing 1 to 12 of 12 entries


Accounting Comparison

WLCG Cloud, Usage

Cloud Resources ▾

MONTH: MAY ▾ YEAR: 2014 ▾

Note: Sorting by multiple columns at the same time can be activated by 'shift' clicking on the column headers which they want to add to the sort. Hovering mouse over the column headers to get descriptions of table columns.

 Search:

Site ▲	Wall time consumer ▾	Wall time provider ▾	CPU time consumer ▾	CPU time provider ▾	Net in consumer ▾	Net in provider ▾	Net out consumer ▾	Net out provider ▾	Memory consumer ▾	Memory provider ▾	Disk consumer ▾	Disk provider ▾
CERN-PROD	103218.75	379976.45	97458.65	195470.63	4536.94	2180	1346.43	1129	3822.75	348160	0.00	3070

Showing 1 to 1 of 1 entries

Cloud Accounting in WLCG

- Sites are the suppliers
 - Site accounting generates invoices
 - For resources used
- Need to monitor the resource usage
 - We trust sites and hence their invoices
 - Comparison can detect issues and inefficiencies
- Job activities in the domain of the VO
 - Measurement of work done
 - i.e. value for money
 - Information not included in cloud accounting
 - Need a common approach to provide information
 - Dashboard?

Comparison to Grid

- Grid accounting = supply side accounting
- No CE or batch system
 - Different information source
 - No per job information available
- Only concerned about resources used
 - Mainly time-based
 - For a flavour
 - A specific composition

Core Metrics

- Time
 - Billed by time per flavour
- Capacity
 - How many were used?
- Power
 - Performance will differ by flavour
- Total computing done
 - $\text{power} \times \text{capacity} \times \text{time}$
- Efficiency
 - How much computing did something useful

Measuring Computing

- Resources provided by flavour
 - How can we compare?
- What benchmarking metrics?
 - And how we obtain them
- Flavour = SLA
- SLA monitoring
 - How do we do this?
- Rating sites
 - Against the SLA
 - Variance

Data Mining Approach

- VOs already have metrics on completed jobs
 - Can they be used to define relative performance?
- Avoids dicussion on bechmarking
 - And how the benchmark compares to the job
- May work for within the VO
 - But what about WLCG?
- Specification for procurement?
 - May not accept a VO specific metric
- Approach currently being investigated

Volunteer Computing


- Should exhaust this approach first
 - Before resorting to commercial providers
- Similar development/operational investment
- vLHC as a common platform
 - Shared community management/support
- Used by
 - ATLAS
 - CMS
 - LHCb

ATLAS@home



Running jobs

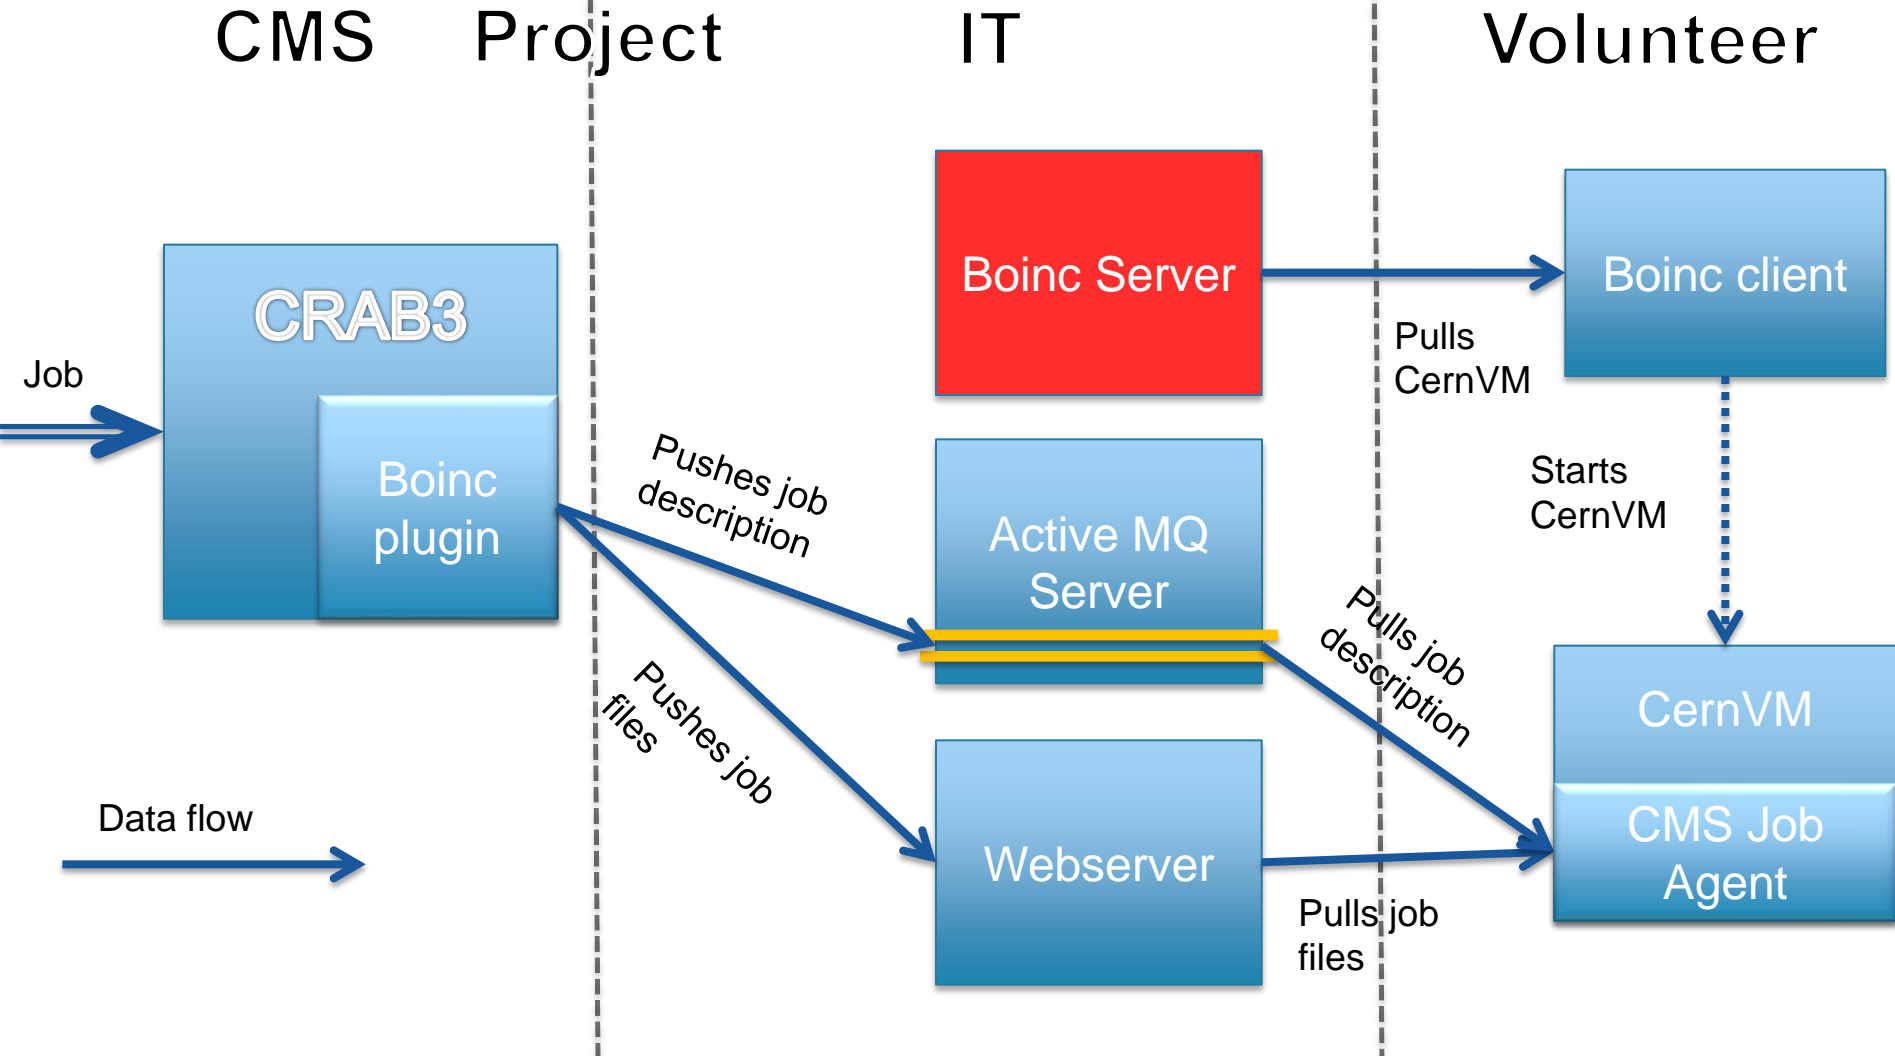
161 Days from Week 13 of 2014 to Week 36 of 2014



Maximum: 4,281 , Minimum: 0.00 , Average: 899.78 , Current: 4,281



CMS@home prototype



Cloud Adoption Document

- Describes the general approach
 - And the status within each experiment
- The first draft to be released within the next few weeks
 - Hopefully will become a focal point for discussion

Summary

- It is all about starting a CernVM image
 - Exploring different options
- How resources are going to be accounted?
 - Counting cores is easy
 - Normalizing for power is hard
- Investigating using job metrics
 - To discover relative performance
- Need to think about procurement
 - What metrics to specify?