


Resource Provisioning Overview

Laurence Field

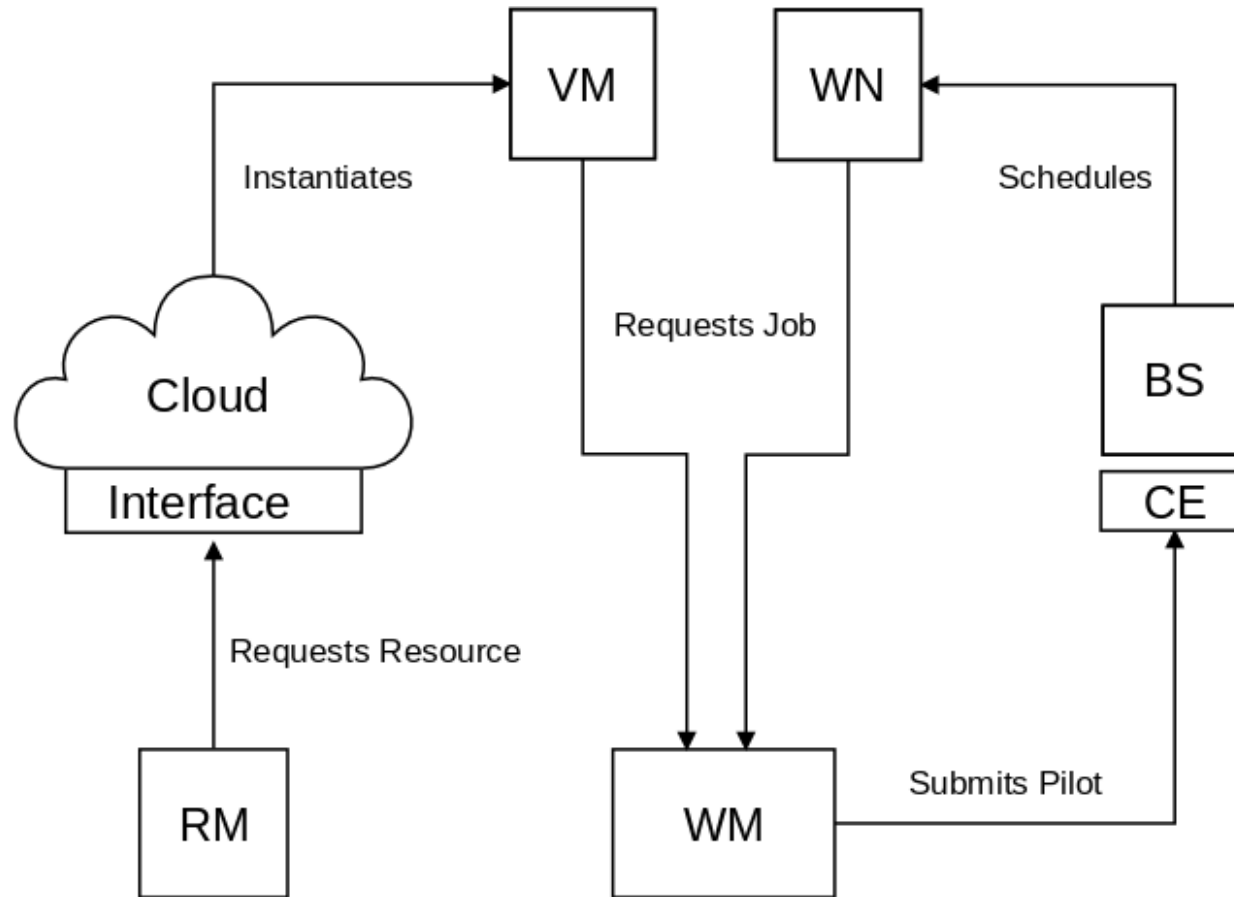
12 April 2015



Overview

- Architectural Model
 - Focus Areas
 - Resource Groups
- Area Details
- Commercial Cloud Activities

From Pilot Jobs to Pilot VMs



Focus Areas

- Image Management
- Capacity Management
- Monitoring
- Accounting
- Pilot Job Framework
- Data Access and Networking
- Quota Management
- Supporting Services

Resource Groups

- WLCG Tier 0/1/2
- HLT Farms
- Commercial Providers
- Volunteer Computing
- HPC
- Parameter space contains 160 possibilities
 - 5 resource groups, 8 areas, 4 VOs
 - Need consolidation of solutions
- The cloud paradigm
 - The adoption of generic solutions
 - Reducing our total cost of ownership

Image Management

- Image provides the job environment
 - PilotJob
 - Pilot VMs
 - Software
 - CVMFS
- Contextualization
 - Balance pre- and post-instantiation operations
 - Simple or Complex
 - Frequency of Updates
 - Data Transferred
- Transient
 - No updates of running machines
 - Destroy (gracefully) and create new instance

CernVM

- The OS via CVMFS
 - HTTP replication of a reference file system
 - Stratum 0
- Why?
 - Because CVMFS is already a requirement
 - Reduces the overhead of distributed image management
 - Manage version control centrally
- CernVM as a common requirement
 - Parameter space reduction
 - 20 possibilities => 1 CernVM
 - 5 resource groups, 4 VOs
 - Availability becomes an infrastructure issue
 - Potentially 20 different contextualizations
 - Responsibility of the VO
- The goal is to start a CernVM-based instance
 - There are exceptions where this is not possible

Capacity Management

- Pilot Jobs => Pilot VMs
- Ensure there are enough resources (capacity)
 - Scale Up/Down
 - Managing the VM life cycle isn't the focus
- Requires a specific component with some intelligence
 - Do I need to start of VM and if so where?
 - Do I need to stop a VM and if so where?
 - Are the VMs that I started OK?
 - Detection of the not ok state maybe non-trivial
- Existing solutions focus on deploying applications in the cloud
 - Difference components, one cloud
 - May not be appropriate for the resource group
 - May manage load balancing and failover
 - Is this a load balancing problem?
 - One configuration, many places, enough instances?

Vacuum Model

- Pilot VMs appear spontaneously out of the vacuum
 - Robust generic approach
 - Can choose whether or not to use the resource
- Pilot Job framework parameter space reduction
 - 20 possibilities => 4 approaches
 - 1 per experiment
- Stimulating resources via framework
 - Optimization
 - Adds complexity
 - Requires additional communication
 - No parameter space reduction
- Area for discussion

Monitoring

- Fabric management
 - The responsibility of the Capacity Manager
 - Basic monitoring is required
- The objective is to triage the machines
 - Invoke a restart operation if it not ok
 - Detection of the not ok state maybe non-trivial
- Other metrics may be of interest
- Spotting dark resources
 - Deployed but not usable
- Can help to identify issues in other systems
 - Discovering inconsistent information through cross-checks
- A common requirement for all
 - Potential for 20 => 1 parameter space reduction
- Jobs monitoring is VO specific
 - Should be no difference with Grid jobs

Provider Accounting

- Commercial suppliers send invoices
 - They have their own (proprietary) accounting systems
- WLCG sites are suppliers
 - Cloud accounting generates monthly “invoices”
 - We trust sites and hence their invoices
- Need to method to record usage
 - To cross-check invoices
 - Detect issues and inefficiencies
- No abstraction of job concept in IaaS layer
 - Job activities are in the domain of the VO
- A job is the VO's measurement of work done
 - i.e. value for money

Consumer Accounting

- Monitoring and recording resource usage
 - What, where, when for VM resources
 - And whatever else is billed for
 - Course granularity acceptable
- Report generation
 - Mirror invoices
 - Use same metrics as the invoice
- Needs a uniform approach for all the VOs
 - Require 20 => 1 parameter space reduction
 - Provide the same information to the budget holder
- Job information is not considered
 - But could potentially be repurposed
- Both monitoring and job information could be used
 - An area for discussion

Other Considerations

- Data access and networking
 - Have so far focus on non-data intensive workloads
- Quota Management
 - Currently mostly fixed limits
 - Leading the partitioning of resources between VOs
 - How can the sharing of resources be implemented?
- Supporting Services
 - What else is required?
 - Eg squid caches in the provider
 - How are these managed and by who?
- Non-Virtualized approaches
 - Instantiation of a pilot job
 - Without CE

Commercial Clouds

- Helix Nebula
 - A public-private partnership
 - Between research organizations and IT industry
- Microsoft Azure Pilot
 - Preliminary discussions with CERN OpenLab
- Amazon
 - BNL RACF for ATLAS and CMS
 - With new Scientific Computing group at AWS
- Deutsche Börse Cloud Exchange AG
 - Beta testing platform
 - Will go live beginning of May
- PICSE
 - Procurement Innovation for Cloud Services in Europe
- European Science Cloud Pilot
 - Pre-Commercial Procurement (PCP) proposal
 - Buyers group public organizations that are members of the WLCG collaboration

Next Steps

- Define and publish contextualizations
 - For each VO
 - For the different resource groups
- WLCG Resource Accounting
 - No significant capacity from Tier 1/2 sites without it
- Commercial Resources
 - Have gained experience for a number of initiatives
 - Who pays?
 - No significant capacity without a budget
 - Need to be ready to integrate commercial resources
- Capacity Management
 - Multiple solutions
 - Continue to experiment
 - Consolidate

Summary

- Resources in many places
 - Not just traditional Tier 0/1/2 sites
- Common approaches
 - To avoid an explosions in our parameter space
- The vacuum model as the generic approach
 - Centered around contextualizing a CernVM
 - To run a pilot job
- Have to consider specializations
 - Commercial clouds
 - Volunteer Computing
 - HPC
- But keep just the special bits special
 - And have common things everywhere else