

Cloudscheduler (v2)

Berghaus F, Casteels K, Driemel C, Ebert M, Leavett-Brown R C, MacDonell D,
Paterson M, Seuster R, Sobie R, Tolkamp S

University of Victoria

Cloudscheduler

Provision CernVM to execute batch workloads on cloud infrastructure

ATLAS - HEP Experiment at CERN

- Cloudscheduler: Integration of cloud resources
- Sim@P1: Opportunistic use of high level trigger resources

Belle-II - HEP Experiment at KEK

CANFar - Canadian Advanced Network For Astronomy Research

- Using cloudscheduler and testing CernVM

Cloudscheduler V2

Version 1 released in 2009 - time for a rewrite

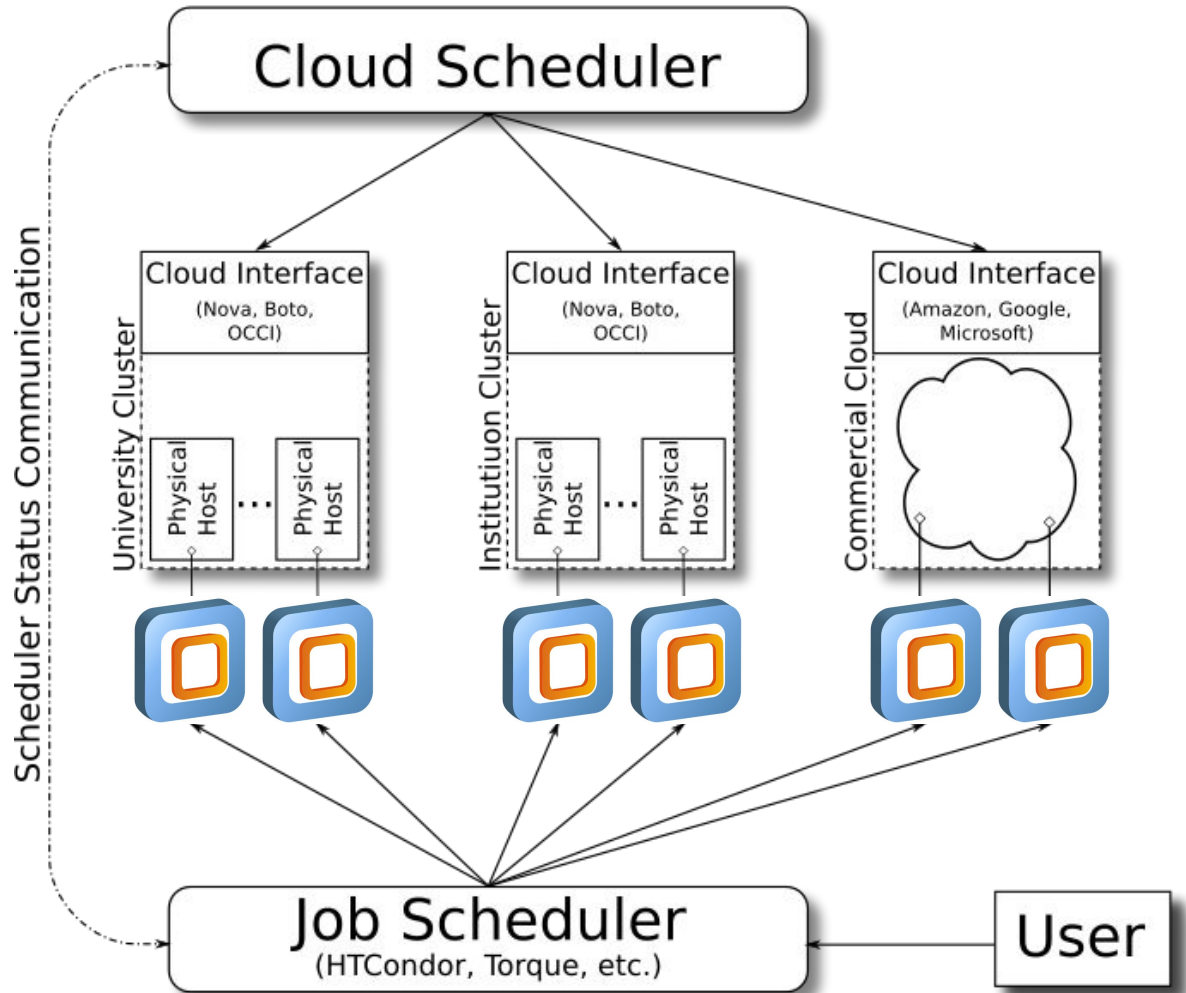
- Python 3, architecture redesign, graphical user interface
- <https://github.com/hep-gc/cloudscheduler>

Version 2




- Driving ATLAS cloud queues
- Future: Belle-II cloud resources, container workflows

Workflow

0. Empty queue, no workers
 - Condor & CS run at UVic
1. User submits job
2. Idle job and no resources
3. Create instance on any cloud with matching resources
 - CernVM batch version
4. Instances connect to Condor
5. Jobs execute
6. No idle jobs in queue
7. CS retires instances
 - Delete on job completion



Reading from services



Instructions to clouds

Architecture

Monitoring

Tables of operational variables

Time series plots!



Image distribution

Image distribution between clouds

Usually CernVM images :)

Lead transition to SL7/CentOS7 with CernVM4 for ATLAS

atlas						
Status	Clouds	Aliases	Defaults	Images	Keys	
Image Filter: <input type="text" value="Search by Image Name"/>						
Toggle Public Images						
+Upload Image						
Submit Changes						
IMAGES	arbutus-a	cc-east-a	cc-west-a	chameleon-a	lrz	otter-a
canarie-demo	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
CentOS 6.6	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
CentOS-6-x86_64-GenericCloud-1711.qcow2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
CentOS-6-x86_64-GenericCloud-20141129	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
CentOS-7-x86_64-GenericCloud-1711.qcow2	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
CentOS-7-x86_64-GenericCloud-1811	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
CentOS-7-x86_64-GenericCloud-1901	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
centos6-bare	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
cernvm	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
cernvm-3.6.5	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
cernvm3	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
cernvm3-micro-2.7-7.hdd	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
cernvm3-micro-2.8-6.hdd	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Configuration

Cloud configuration is info from resources file


Adding support for more cloud infrastructures and configurations

Running already on many clouds (OpenStack and EC2)

The screenshot displays the 'Configuration' interface for the 'Atlas' system, specifically the 'Clouds' tab. The interface is organized into several sections:

- Navigation:** A top bar contains tabs for 'atlas', 'Status', 'Clouds' (active), 'Aliases', 'Defaults', 'Images', and 'Keys'.
- Cloud List:** A sidebar on the left lists cloud configurations: 'arbutus-a', 'cc-east-a', 'cc-west-a', 'chameleon-a', 'lrz', and 'otter-a' (highlighted in red). Below the list are buttons for 'Settings', 'Metadata', and 'Exclusions', followed by a '+' icon.
- Configuration Fields:** The main area shows settings for the selected 'otter-a' cloud:
 - Enabled:** A checked checkbox.
 - URL:** A text input field containing 'https://otter.heprc.uvic.ca:'.
 - Project:** A text input field containing 'HEP'.
 - Username:** A text input field containing 'Atlas'.
 - Password:** A text input field with a placeholder 'Update password'.
 - CA certificate:** A text input field containing '/etc/ssl/certs/CABundle.cr'.
 - Region:** A text input field containing 'Victoria'.
 - User domain name:** A text input field containing 'Default'.
 - Project domain name:** A text input field containing 'default'.
 - Spot Price:** A text input field containing '-1'.
- VM Settings:** A separate section on the right contains:
 - Cloud type:** A dropdown menu set to 'openstack'.
 - VM Keyname:** An empty text input field.
 - VM Network:** A dropdown menu set to 'private'.
 - VM Image:** An empty text input field.
 - VM Flavor:** A dropdown menu set to 'b8'.
 - VM Keep Alive:** A text input field containing '-1'.
- Resource Limits:** A section at the bottom right features:
 - Cores Softmax:** A text input field containing '400'.
 - Cores:** A slider control with a value of '400' out of a maximum of '500'.
 - RAM:** A slider control with a value of '-1' out of a maximum of '1023983 KB'.
- Actions:** An 'Update Cloud' button is located at the bottom center of the configuration area.

Running slots by PanDA queue



	min	max	avg	current
CA-IAAS-T3-K8S	0	163	95	145
CA-IAAS-T3_UCORE	0	1.428 K	477	554
CERN-EXTENSION_HARVESTER	0	838	668	3
HEPHY-UIBK	0	1.800 K	1.335 K	1.715 K
LRZ-LMU_CLOUD	0	413	149	188
UKI-SCOTGRID-ECDF_CLOUD	0	373	172	343


Cloudscheduler in ATLAS

Used in commissioning & integration:

- Singularity, containers, pilot2, HTTP/WebDAV,

...

Normalized CPU usage by Country 14 Days from 2019-05-19 to 2019-06-02

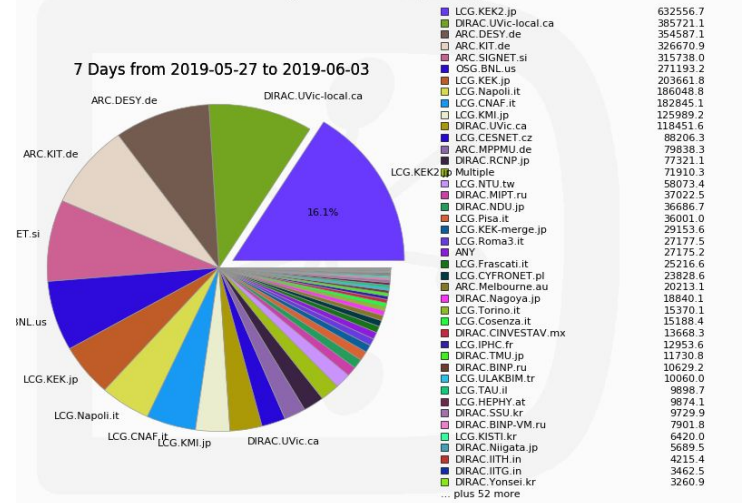


Max: 276, Min: 37.8, Average: 189, Current: 37.8

JP	26.0%	TW	2.7%	MX	0.7%	CN	0.0%
CA	19.7%	RU	2.7%	KR	0.5%	ANY	0.0%
DE	18.7%	CZ	1.4%	IL	0.4%	MULTIPLE	0.0%
SI	11.6%	PL	0.9%	IN	0.3%		
IT	9.6%	FR	0.8%	AT	0.2%		
US	2.8%	AU	0.8%	TR	0.2%		

Generated on 2019-06-03 14:07:34 UTC

Average Number of Jobs by Site



Generated on 2019-06-03 17:30:37 UTC

Cloudscheduler in Belle-II

Provides Canadian computing pledge
Uses Dynafed as data access layer (gfalfs)

Other notable CernVM & CVMFS related work

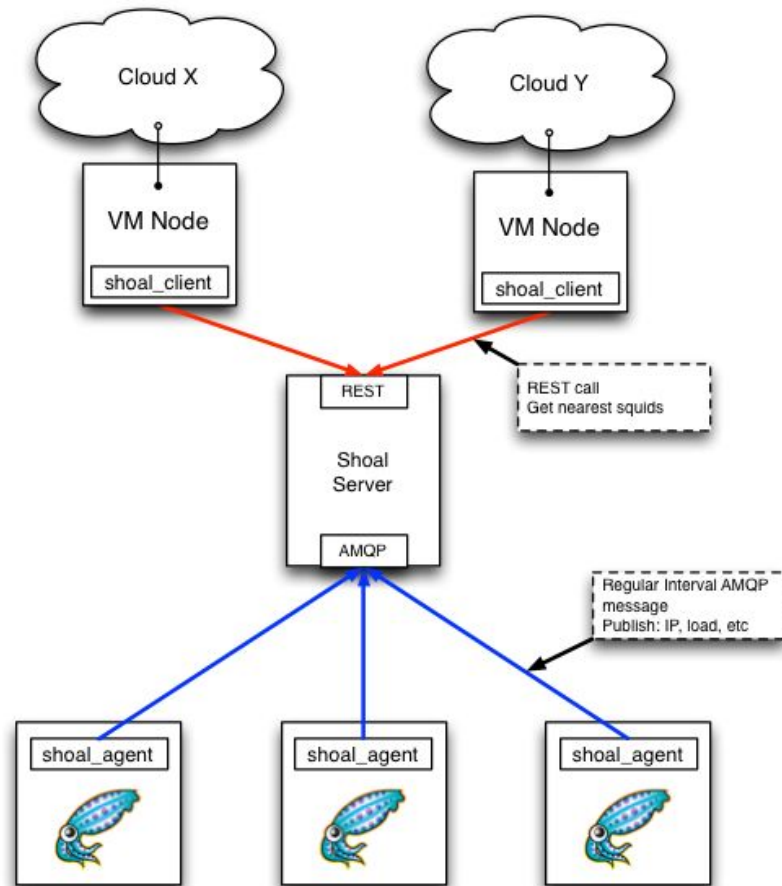
Shoal

(Squid) proxy discovery

shoal.heprc.uvic.ca

Used for CVMFS and frontier on clouds

1. Agent on cache advertises to server
2. Server performs function test on cache
3. Server registers cache
4. Client requests cache(s)
5. Server calculates best match(es)
 - a. ACL, GeoIP, and load
6. Server responds to client



Sim@P1

Opportunistic use of ATLAS HLT farm

O(100k) cores

Virtualization provides isolation

Boot with libvirt tools

Configuration with config disk

Three (3) dedicated squids

Boots 2,500 instances in 2 minutes



Conclusion

Cloudscheduler orchestrates CernVM instances

CernVM provides stable system that allows us to test and develop

Container testing ongoing

- Pilot calls singularity
- Kubernetes manages containers running Pilot