

1 **High-throughput cloud computing with the**  
2 **cloudscheduler VM provisioning service**

3 **F. Berghaus, K. Casteels, C. Driemel, M. Ebert, F. F. Galindo\***,  
4 **C. Leavett-Brown, D. MacDonell, M. Paterson, R. Seuster,**  
5 **R. J. Sobie, S. Tolkamp, J. Weldon**

6  
7 Received: date / Accepted: date

8 **Abstract** We describe a high-throughput computing system for running jobs on public and private computing  
9 clouds using the HTCondor job scheduler and the cloudscheduler VM provisioning service. The distributed  
10 cloud computing system is designed to simultaneously use dedicated and opportunistic cloud resources at local  
11 and remote locations. It has been used for large scale production particle physics workloads for many years us-  
12 ing thousands of cores on three continents. A decade after its initial design and implementation, cloudscheduler  
13 has been modernized to take advantage of new software designs, improved operating system capabilities and  
14 support packages. The updated cloudscheduler is more resilient and scalable, with expanded capabilities. We  
15 present an overview of the original design and then describe the new version of the distributed compute cloud  
16 system. We conclude with a review of the current status and future plans.

17 **Keywords** particle physics · cloud computing

## 1 Introduction

In the field of particle physics, clouds are increasingly used to process and analyze data [1]. It is possible that commercial clouds might satisfy the computational requirements of the next generation of global research projects; however, it is likely that dedicated storage facilities will be required to store and preserve the research data. One scenario may be a hybrid solution of dedicated and opportunistic, private and commercial compute resources, linked by high-speed networks to a distributed set of dedicated storage repositories.

This paper describes such a hybrid solution for running high-throughput computing workloads on compute clouds, irrespective of the underlying cloud software, location or its ownership. We call our design a *distributed compute cloud* where the resources are an aggregate of compute clouds hidden from the user. The distributed cloud can be integrated with existing storage systems or federations to provide full access to the research data. Although the focus of the distributed cloud is to deliver resources for particle physics application, it is also used for astronomy and can be used by researchers in other fields.

The distributed compute cloud uses cloudscheduler for VM provisioning and scheduling, and HTCondor for job scheduling. The two packages are designed for a dynamic environment where the resources on a cloud, or even the cloud itself, appear or disappear. Briefly, cloudscheduler reviews the HTCondor job queue and cloud resources to determine whether there are clouds that can start VMs that meet the job requirements. If a match is found, then cloudscheduler requests that the appropriate VM image be booted on the cloud. Once the VM is instantiated, it joins the HTCondor pool and the user job is sent to the running VM instance.

The original design of cloudscheduler was conceived in 2009 [2] and is based on ideas discussed in a paper describing “sky computing” [3]. The distributed compute cloud has run many millions of jobs on three continents for the ATLAS experiment at the CERN Laboratory in Geneva, Switzerland [4] and the Belle II experiment at the KEK Laboratory in Tsukuba, Japan [5].

Many of the custom and external software components have undergone significant change since the first version of cloudscheduler. Further, the years of production operation have provided insight on how to simplify and improve the system (see section 2.3). In 2018, cloudscheduler was changed from a single Python code framework into a software platform with expanded functionality using current software technologies and practises. The new design is more robust, error tolerant and scalable.

There are other strategies and software for utilizing clouds in particle physics. These include Vcycle, VMDIRAC and an extension of the HTCondor job scheduler. In Vcycle, the resource provider creates VMs for the experiments that draw jobs from the experiments’ central queue of tasks [6]. VMDIRAC is a module for the DIRAC workload management system [7] that can start VMs on clouds [8,9]. HTCondor can also be used to start VMs on Openstack clouds [10]. The distributed cloud computing system using cloudscheduler provides an infrastructure that can run workloads for any field or research without the need for application experts at each site. It is not dependent on project-specific workload management systems but can be integrated with them. Further, it can run workloads on all cloud types while the users or experiments only need to know about the familiar batch system interface to which they submit jobs.

In this paper, we give an overview of the original design of the distributed cloud system using Cloudscheduler Version 1 (CSV1) and motivation for a new version in Section 2. Section 3 describes Cloudscheduler Version 2 (CSV2). Throughout the paper, we highlight the external components used in both versions of the distributed compute cloud that are critical to the system.

## 2 Cloudscheduler V1 distributed compute cloud

### 2.1 Overview

We briefly describe CSV1 to give some context and motivation for the new version. The architecture of the distributed compute cloud using CSV1 has been described in a number of papers [11,12] and is shown in Fig. 1. The key components are the HTCondor job scheduler, the cloudscheduler VM provisioning service and the compute clouds. CSV1 provides API support for Openstack, Open Nebula, Amazon EC2, Microsoft Azure and Google GCE clouds.

HTCondor was selected as the job scheduler as it was designed to be a cycle scavenger [13], making it an ideal fit for a dynamic cloud environment. The user or workload management system submits a job to the HTCondor job scheduler, specifying the job requirements (e.g. number of cores, memory and disk space) in the Job Description Language (JDL) file. Multi-core VMs are booted on the clouds whether a single job requires

one or all of the cores in the VM as there are benefits such as shared disk caches, fewer VM instances and it helps reduce the fragmentation of the resources on the clouds. The HTCondor client on the VM starts a partitionable slot during the contextualization process, which is subdivided depending on the resource requirements of the job.

CSV1 examines the list of pending jobs in the HTCondor queue and searches for a cloud with free resources that meets the job requirements as specified in the JDL file or as specified by the system-wide job defaults. If there is a cloud that meets those requirements, then CSV1 sends a request to the cloud to boot a VM. CSV1 only requests one VM boot per cycle every minute (both configurable parameters) on one cloud in the list and on the subsequent cycle, it would request a VM on the next cloud in the list. The VM image and VM flavour are specified in the JDL file or in a configuration file. Once the VM is booted and contextualized, it joins the HTCondor pool and queued job(s) can start on this VM.

The CSV1 workflow is executed approximately once per minute. Decisions on VM provisioning are based on the information about the jobs and clouds. CSV1 retains the state information in memory and can write it to disk using the Python pickle module. The pickle file makes it possible to restart CSV1 if there are minor issues or there is a simple code change. More significant changes or outages require the entire system to be shutdown.

CSV1 can request the start, retirement or immediate destruction of a VM. The start and destroy commands are directly issued via the cloud API. If there are no jobs in the HTCondor queue, then a running VM may be retired. A retire request is issued to the HTCondor client on the VM and the client deregisters from the HTCondor pool. The retire request issues the “condor.off” command to both the HTCondor startd and master daemons on the VM instance. The HTCondor slots on this machine are then listed as being in a “retired” state. The jobs running on the VM are allowed to finish and then the VM is destroyed.

HTCondor and CSV1 are managed via their respective command line interfaces (CLIs). The Linux root user can configure the system, add or remove clouds and edit the files used for the contextualization of a VM. The normal Linux users can query the status of jobs and VMs using the HTCondor and CSV1 CLIs, respectively. CSV1 does not have a GUI for managing the system, although there is a monitoring web page.

## 2.2 Performance and status

The distributed compute cloud is integrated into the WLCG grid infrastructure [14] and has run production workloads for many years using clouds in North America, Europe and Australia. The majority of clouds use the Openstack software, though Open Nebula, Amazon EC2, Google GCE and Microsoft Azure clouds have also been used.

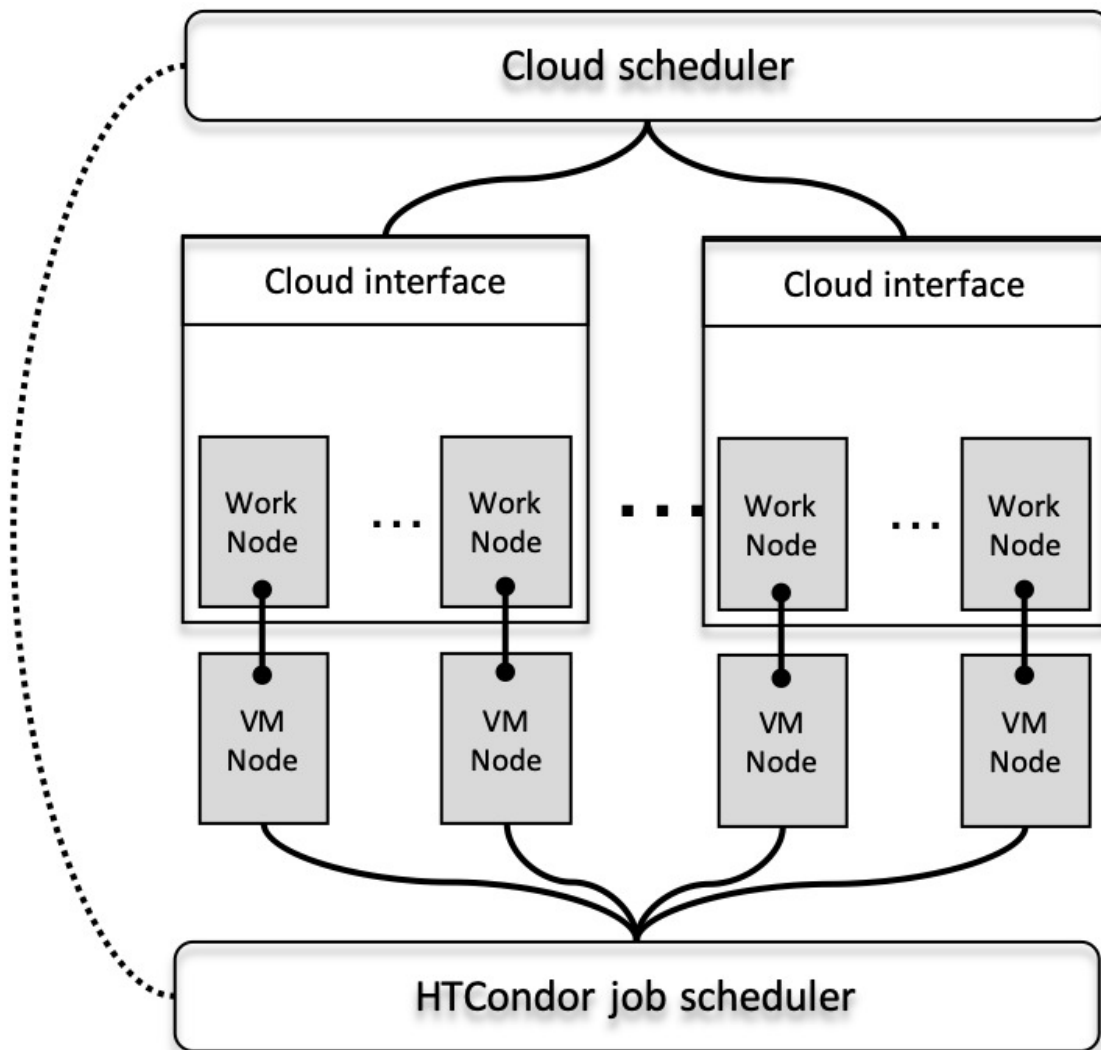
We ran two instances of the CSV1 system in North America for ATLAS and Belle II, respectively, and another CSV1 in Europe for ATLAS (each with a separate HTCondor instance). The HTCondor instance dedicated to the ATLAS experiment is linked to the PanDA workload management system [15]. ATLAS typically uses a few hundred thousand cores at a given moment, and the distributed compute cloud contributes approximately 1% of the resources, comparable to the other Tier-2 compute centres operated in Canada. ATLAS (and Belle II) submit *pilot jobs* that sets up the software environment and contacts the central system for payload jobs. A pilot job can run more than one payload job depending on its configuration.

The Belle II experiment uses the DIRAC workload management system [7]. We operate three DIRAC Site Directors in Victoria that submit pilot jobs to the Belle II HTCondor instance if there are queued jobs in the central DIRAC server at KEK. The distributed compute cloud has provided 13% of the total workload of the experiment since January 2018. All of the Canadian resources for Belle II will continue to be provided by the distributed compute cloud.

## 2.3 Motivation for cloudscheduler redevelopment

Cloudscheduler was initially developed in 2009 and evolved with the cloud technologies but it has become difficult to maintain. The CSV1 code is written in Python 2, which will no longer be supported after January 2020. CSV1 also contains redundant code written to support operating system features and cloud software that no longer exist. CSV1 was designed as a multi-user system; however, it never functioned well as a multi-user system as it would rapidly start and stop VMs in attempt to balance resources between competing projects.

CSV1 runs on a single node and uses in-memory data structures for the state information; making it prone to single errors and failure. An unscheduled outage or system crash would result in a loss of all running VMs



**Fig. 1** Overview of the cloudscheduler Version 1 and HTCondor distributed compute cloud. A user or workload management system submits application jobs to the HTCondor job scheduler. Cloudscheduler reviews the job queue and the resources on the compute clouds. If there is a cloud with resources that meet the requirements of the job, then cloudscheduler issues a command to boot the user-specified VM. The instantiated VM registers with HTCondor and the job is submitted to the VM.

119 and require a manual intervention on each cloud to remove the residual VMs (a time consuming and error prone  
120 task).

121 As mentioned in the previous section, there were separate instances of CSV1 for the experiments. While  
122 the separate CSV1 instances worked well, it meant one could not shift idle resources between projects. A single  
123 instance of cloudscheduler would also simplify the operations though it would need to be able to use one or  
124 more instance of HTCondor. The ATLAS and Belle II experiments use different authentication in HTCondor,  
125 and some sites do not allow their VMs to connect to a remote HTCondor pool.

126 It became apparent that a new version was necessary to update the software, simplify the operations, extend  
127 the current capabilities and add new functionality.

### 128 3 Cloudscheduler V2 distributed compute cloud

#### 129 3.1 Overview

130 CSV2 is a framework of component processes written in Python 3 and built around the MariaDB (SQL) re-  
131 lational database [16]. There are processes for control, data gathering (pollers), user interfaces and for VM  
132 provisioning. The MariaDB ensures consistent state information, which is essential for the reliability of CSV2  
133 (see Fig. 2).

134 The database allows data retrieval from multiple sources, and correlates information between different  
135 subsystems to obtain different views of the system data. The data is organized in tables, with rows representing  
136 objects and columns representing the attributes of an object. There are tables for global and local configuration  
137 parameters. CSV2 also generates ephemeral data in dedicated tables for objects like jobs or virtual machines.  
138 The database also adds authorization features that protects the integrity of the data while allowing both local  
139 and remote access.

140 The database also ensures recoverability (assuming a regular database backup strategy). Should a catas-  
141 trophic system loss occur, e.g. through hardware failure, restoration of the configuration data from the latest  
142 database backup will allow CSV2 to restart and return to full operation. Within one cycle, the data pollers (de-  
143 scribed below) will retrieve information from all configured subsystems (clouds and HTCondor job schedulers),  
144 and by updating the ephemeral data, it will allow the database to accurately reflect the current state once again.

145 CSV2 introduces a RESTful web User Interface (UI) for administration, control, and detailed monitoring of  
146 the clouds and jobs. In contrast to CSV1, the CSV2 UI can be accessed through either a graphical user interface  
147 (GUI) using web browsers or a command line interface (CLI) client (with extensive man pages); the GUI and  
148 CLI have nearly the same functional capabilities. Different to CSV1, the CLI can be used from any computer  
149 and the user of the CLI does not need to have ssh access to or a Linux account on the machine where CSV2  
150 runs.

151 For all clients, the connections are authenticated either by x509 proxy certificates or by usernames and  
152 passwords. In Fig. 3, an example of the graphical status display is presented.

153 CSV2 has *pollers* that fill the MariaDB with information from the clouds and HTCondor. New entries are  
154 added to multiple tables in the database for each new job and VM. The state information of existing entries are  
155 updated by the *pollers*, and obsolete information (for completed jobs/VMs) is removed.

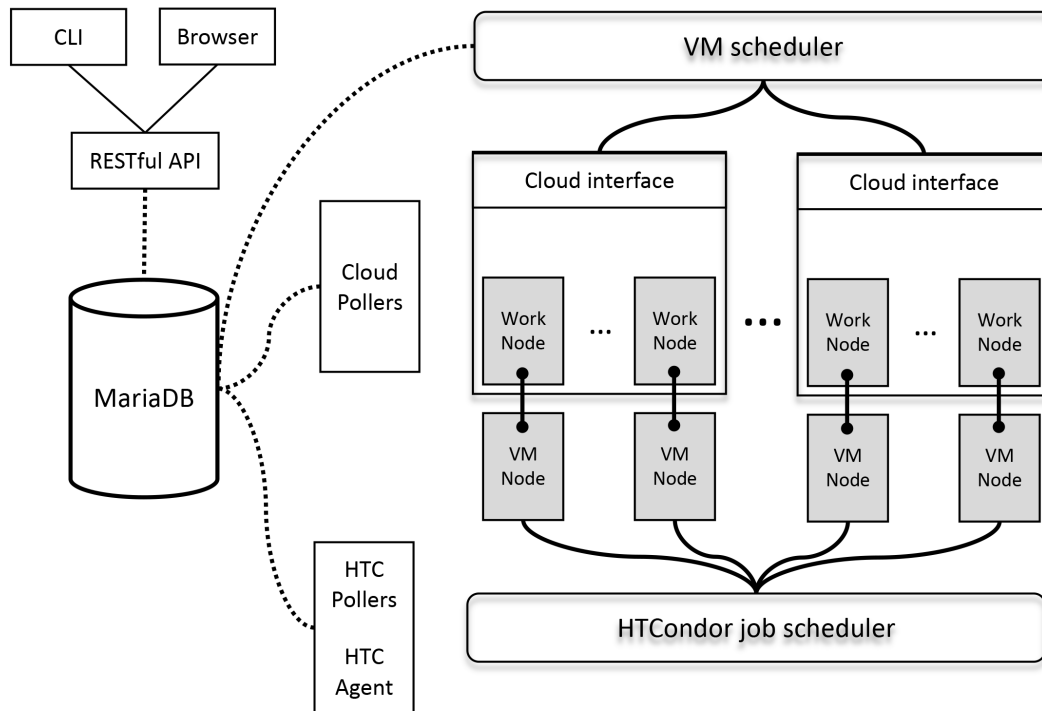
156 Fig. 2 shows a single *cloud poller*; however, there are separate *cloud pollers* for each cloud-type<sup>1</sup> (e.g.  
157 one poller can support multiple Openstack clouds). Fig. 2 also shows pollers for HTCondor (labelled as *HTC*  
158 *pollers*). There is a *HTC machine poller* that gathers information on the HTCondor machines and a *HTC job*  
159 *poller* that gathers information on the jobs in the HTCondor queue. The pollers consist of multiple tasks that  
160 run at different frequencies. For example, the *cloud poller* updates the MariaDB with information on the VMs  
161 every 60 seconds whereas it updates the list of VM images every 5 minutes. Typically, the dynamic information  
162 is updated every minute and the relatively static information on a longer timescale (between 5 minutes and one  
163 day)<sup>2</sup>.

164 Although Fig. 2 shows only one HTCondor instance, CSV2 can support multiple instances. For example,  
165 separate HTCondor instances are used by ATLAS and Belle II experiments, in part, due to differing security  
166 requirements. All HTCondor instances are supported by a single set of HTCondor pollers. However, each  
167 instance of HTCondor requires a *HTC Agent* to provide the correct security context when issuing requests to  
168 the HTCondor client on the VM (e.g. requesting a VM to de-register from the HTCondor pool).

169 CSV2 requires inter-process communication and signaling. For example, the User Interface (UI) wakes the  
170 cloud poller when a cloud is added, so that cloud data is updated immediately or the machine poller signals  
171 the HTC Agent to retire a VM. CSV2 uses AMQP, a reliable message delivery protocol, to signal processes to  
172 wake when there is important information that will impact the outcome. For example, the *cloud poller* looks for  
173 new clouds or VM images every 5 minutes. The AMQP message from the UI, signals the *cloud poller* to start a  
174 new cycle. This solution is better than relying on the pollers running on an overly frequent cycle. It is expected  
175 that AMQP will have an expanded role in further updates of the system.

<sup>1</sup> Currently there is API support for Openstack and Amazon EC2 clouds (support for other cloud APIs will be added on demand)

<sup>2</sup> Note that the numbers given for cycle times, limits and thresholds are configurable parameters. Further, when we state that a poller or process is run every  $N$  seconds, this means that the process sleeps for  $N$  seconds after the last cycle before restarting its tasks.



**Fig. 2** Overview of the cloudscheduler Version 2 and HTCondor distributed compute cloud. The primary change is the introduction of the MariaDB for keeping track of the state of the system and “pollers” to gather information from the clouds and the job scheduler. The information in the database is used for scheduling, management and monitoring.

**Table 1** The possible states of a VM in the CSV2 distributed compute cloud system

VM state	Description
<i>starting</i>	VM is booting/contextualizing
<i>unregistered</i>	VM is running and has not registered in HTCondor pool
<i>idle</i>	VM is running, registered in HTCondor pool and not running jobs
<i>running</i>	VM is running, registered in HTCondor pool and running jobs
<i>retiring</i>	VM is running, retired in HTCondor pool and will complete running jobs
<i>manual</i>	VM is flagged as being manually used and will be ignored by the <i>VM Scheduler</i>
<i>error</i>	VM in error state according to the cloud information

### 176 3.2 Workflow

177 The VM Scheduler runs every 10 seconds and retrieves virtual tables (views) from the MariaDB to determine  
 178 its actions. A view can retrieve and relate data from multiple tables, perform calculations, sort, group, select,  
 179 concatenate, resort and regroup to derive the information required. One of views determines the state of the  
 180 VMs (see Table 1). The primary tables required by the scheduler are the job queue and the compute resources.

181 We briefly describe the steps to create the job queue view. Jobs are grouped according to their resource  
 182 requirements where resources such as cores, disk, and RAM are considered. For example, both ATLAS and  
 183 Belle II have simulation, reconstruction and analysis jobs requiring either 1 or 8 cores and varying amounts of  
 184 RAM and disk. Jobs of similar requirements are grouped and counted (called a “job-group”).

185 Next, the view finds the VM flavours for each cloud that can satisfy the job requirements of the job-group.  
 186 For each cloud, the view assembles a list of (Cloud name:VM flavour) pairs with only one pair per cloud. The  
 187 VM flavour is specified by the user or retrieved from the database. Most clouds have many VM flavours and  
 188 when the VM flavour is obtained from the database, then the one with the fewest cores and/or the least amount  
 189 of memory is selected. If a cloud has no VM flavour that meets the requirements, then that cloud is not included  
 190 in the list of VM flavors (and will not be selected to boot VMs for this job-group).

191 The result of the view is the table shown below, listing the group name, the number of idle jobs, and a list of  
 192 (Cloud name:VM flavour) for each unique job type (only a subset of the columns in the table is shown). Each  
 193 row corresponds to a unique job type. The first row is for ATLAS 8-core jobs, the second row is for Belle II  
 194 production (low I/O) jobs and the third row is for Belle II analysis (high I/O) jobs. The latter group requires a  
 195 local Belle II storage element, which is set to run on different set of clouds.

Group	Idle jobs	Cloud name : VM flavour
atlas	86	arbutus-nf:c16-60gb-392, arbutus:c8-30gb-186, cc-east:c8-30gb-430, chameleon:m1.xlarge, otter:b8
belle	1	arbutus:c8-30gb-186, cc-east:c4-15gb-205, desy:m1.xlarge, ecdf-b:m1.xlarge, otter:b8
belle	91	arbutus:c8-30gb-186, otter:b8

197 The list of (Cloud name:VM flavour) in the above table is ordered by the priority of the cloud (called the  
 198 cloud-priority). The cloud-priority can, for example, give lower priority to commercial clouds. Clouds that have  
 199 the same priority are ordered alphabetically based on their cloud name.

200 Prior to reviewing whether the clouds in the job-group have free resources to boot new VMs, the system is  
 201 checked for VMs in a *starting* or *unregistered* states (see Table 1), or if there are *running* VMs that are not fully  
 202 utilized (e.g. empty or idle job slots). This indicates there is capacity in the system for the idle jobs and no new  
 203 VMs will be booted.

204 The compute resources view is used to determine how many VMs with a specific flavour can be started on  
 205 one cloud. The following table shows an example of a view that returns the group name, the (Cloud name:VM  
 206 flavour), and the number of Available VM slots that can be started on this cloud. In this example, the arbutus  
 207 cloud has the capacity to start 109 VMs with the c8-30gb-186 flavour for ATLAS and 184 VMs with the c8-  
 208 30gb-186 flavour for Belle II. The table was selected to show only the results of a query for a VM flavour of  
 209 arbutus:c8-30gb-186 (an unrestricted query returns over 100 rows).

Group	Cloud name : VM flavour	Available VM slots
atlas	arbutus:c8-30gb-186	109
belle	arbutus:c8-30gb-186	184

211 CSV2 makes it possible for one experiment to opportunistically utilize the resources of another experiment  
 212 whose resources are idle. The number of available VM slots for each cloud and VM flavour (as shown in  
 213 the above table) takes into account opportunistic resources. In the simplest configuration, the primary group  
 214 (experiment) has the priority use of a cloud, and the opportunistic group can also use that cloud if there are  
 215 idle resources. The primary group has a quota of compute cores equal to the full allocation on the cloud, and  
 216 the opportunistic group has an opportunistic quota of compute cores that is less than the full allocation. If  
 217 the primary group does not use any resources, the opportunistic group can use the cloud up to their opportunistic  
 218 quota. If the primary group begins submitting jobs, then the VMs of the opportunistic group will be retired.  
 219 Eventually there will be no jobs from the opportunistic group. The quotas of the primary and opportunistic  
 220 groups are used when determining the compute resources view. This strategy has significantly improved the  
 221 utilization of the clouds.

222 The compute resources view provides the *VM scheduler* with the information it requires to boot new VMs.  
 223 Prior to issuing VM boot requests, a number of conditions are tested to avoid starting too many VMs. For  
 224 example, if there are five or more VMs in a *starting* or *unregistered* state on a cloud, then no VMs will be  
 225 started on that cloud for any group. In addition, if there are more than ten *idle* VMs within the clouds associated  
 226 to a unique job type, then no VMs are started for that job type group.

227 If the (Cloud name:VM flavour) in the compute-view table shows available VM slots (see above table), then  
 228 the *VM scheduler* will issue up to five VM boot requests on each cloud in the group.

229 At the end of the boot process, the VMs are contextualized with the configuration required by CSV2 (e.g.  
 230 HTCondor), the particle physics software (e.g. CVMFS file system) and the experiment software. We discuss  
 231 this in further detail in 3.4.

232 Once the VM has completed its contextualization, it is registered as a HTCondor machine. The *HTC ma-*  
 233 *chine poller* periodically retrieves the ClassAds of all the HTCondor machines via the HTCondor Python API  
 234 and stores its ClassAd in the MariaDB. If there is no ClassAd for the VM, then a query of the MariaDB would  
 235 show the VM to be in an *unregistered* state. The registration of the VM with the HTCondor pool happens during  
 236 the start of the HTCondor clients on the VM after it is fully contextualized. A VM that is in an *unregistered*  
 237 state that exceeds a “come alive” time (currently 2400 seconds) or fails to start any job within a “job alive” time  
 238 (currently 300 seconds) is usually problematic and is terminated.

239 If a VM remains idle for an extended period of time after completing one or more jobs, then either the  
 240 job queue is empty, there are no more queued jobs that can run in the VM, or the VM developed problems  
 241 communicating with the HTCondor server. In this situation, the *VM scheduler* sets the retire-flag in the database  
 242 entry of the VM. We have observed that keeping *idle* VMs alive for 30-60 minutes is optimal (defined by a  
 243 “keep alive” time) due to the sporadic nature of the workload management systems of the experiments. Both  
 244 the ATLAS and Belle II workload management systems keep track of the idle or queued and running jobs, and  
 245 use those numbers to submit additional jobs.

246 If the retire-flag of the VM in the MariaDB is set, then the *HTC machine poller* sends retire messages  
 247 via AMQP to the *HTC Agent*. On receiving a retire message, the *HTC Agent* issues a “condor off” command  
 248 directly to the HTCondor master daemon running on the VM, causing all the partitions to be marked “retiring”,  
 249 preventing them from accepting new jobs. However, jobs currently running on the VM are allowed to complete.”

250 When a VM is in a *retiring* state and it is not running any jobs, then the *HTC machine poller* sends a request  
 251 (via the cloud API) to the cloud to shut down that VM instance.

252 A VM can also be retired if there is a change in the quota of a cloud (e.g. cores, RAM) and the current  
 253 usage now exceeds the new quota. Depending on the type of change in the quota, then either the user interface  
 254 or the cloud poller will set the retire flag on a collection of VMs (based on information such as the age of the  
 255 VM) needed to rebalance the system.

### 256 3.3 Configuration

257 CSV2 introduces users and groups to the distributed compute cloud (stored in the MariaDB). Users can be  
 258 assigned different access privileges and can be members of multiple groups. Groups are defined, in our case, to  
 259 be particle physics experiments or a development and testing area<sup>3</sup>. Each CSV2 group adds the compute clouds  
 260 to its configuration and users in the group have the ability to add and manage the clouds.

261 The VM images and SSH keys<sup>4</sup> can be uploaded by logging in directly to the cloud (e.g. the Openstack or  
 262 Amazon dashboards) or via the CSV2 GUI. One can define a default VM image and default SSH keys that will  
 263 be automatically distributed to all Openstack clouds. CSV2 keeps track of the VM images in the MariaDB. As  
 264 the list of VM images can change, the *cloud poller* periodically queries the clouds for the list of available VM  
 265 images as well as VM flavours and cloud quotas to keep the MariaDB up to date.

266 VM instances are sized and instantiated using flavours (also known as “instance types” in EC2 parlance).  
 267 Users within a group may also define default flavours on a cloud by cloud basis and/or a group default flavour.  
 268 Where both group and cloud default flavours are defined, the cloud default flavour will take precedence over the  
 269 group default flavour for any particular cloud. If a job specifies no flavour, and no default flavours are found,  
 270 then CSV2 will select one that uses the least resources and satisfies the specified job requirements (cores, disk,  
 271 and RAM).

272 The GUI is accessed through a web browser and can also generate time series plots for each variable  
 273 (stored in an InfluxDB). An example of a time series plot is shown in Fig. 4 where the number of running jobs  
 274 for ATLAS and Belle II is plotted for the month of July 2019. For most of 2019, ATLAS provided a sustained  
 275 workload of longer running production jobs (typically 6-12 hours) while Belle II has periodic bursts of very  
 276 short running analysis jobs (5-10 minutes). The time series plot highlights the difference in workload between  
 277 the two projects for June-July 2019. Belle II is still at an early stage of its project lifetime and just starting  
 278 to record particle collision data; we expect more production jobs and longer analysis jobs with the increasing  
 279 volume of data.

280 There are other configuration settings for the management of VMs with default parameters (for example, the  
 281 cycle times of the pollers and VM scheduling algorithm parameters). We refer the reader to the documentation  
 282 located at [cloudscheduler.readthedocs.io](http://cloudscheduler.readthedocs.io)<sup>5</sup>.

<sup>3</sup> Openstack has projects and users. Previously, Openstack used tenants before changing to projects. As CSV2 uses other types of clouds, it was decided to identify ensembles of users as groups to distinguish it from Openstack projects.

<sup>4</sup> The SSH keys allow the user to log into a VM instance for debugging purposes.

<sup>5</sup> The documentation is work in progress

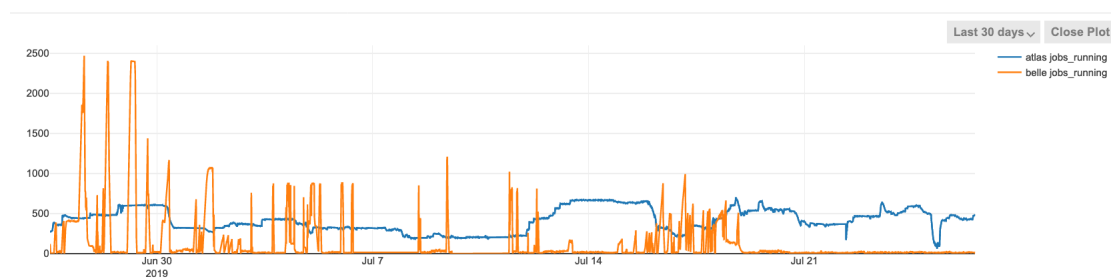


Group	Jobs	Idle	Running	Completed	Held	Other	Foreign
atlas	875	372	483	20	0	0	0
belle	31	4	18	0	9	0	0

Group	Clouds	VMs	Starting	Unreg.	Idle	Running	Retiring	Manual	Error	Native Cores Used	Cores Limit	RAM
atlas	arbutus	309	0	0	0	309	0	0	0	2472	2500	<div style="width: 98%;"></div>
atlas	cc-east	20	0	0	0	20	0	0	0	160	160	<div style="width: 100%;"></div>
atlas	chameleon	12	0	0	0	12	0	0	0	96	96	<div style="width: 100%;"></div>
atlas	otter	50	0	0	0	47	0	0	3	400	400	<div style="width: 100%;"></div>
belle	arbutus	3	0	1	0	2	0	0	0	12	2500	<div style="width: 0.5%;"></div>
belle	cc-east	0	0	0	0	0	0	0	0	0	100	<div style="width: 0%;"></div>
belle	desy	2	0	0	0	2	0	0	0	16	70	<div style="width: 23%;"></div>
belle	ecdf-b	0	0	0	0	0	0	0	0	0	64	<div style="width: 0%;"></div>
<b>Totals</b>		<b>396</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>392</b>	<b>0</b>	<b>0</b>	<b>3</b>	<b>3156</b>	<b>5890</b>	<div style="width: 53.6%;"></div>

**Fig. 3** Snapshot of the CSV2 GUI showing the number of jobs for the ATLAS and Belle II experiments in the upper table. The second table shows the different clouds running for the two experiments and the states of the VM instances. At the time of screenshot, the Belle II experiment had few jobs running and the resources were mainly used by ATLAS. Only part of the GUI is shown.



**Fig. 4** Shows the number of running ATLAS jobs (blue) and Belle II jobs (orange). ATLAS was running production jobs requiring 4 or 8 cores whereas Belle II was periodically running very short analysis jobs during the 30-day period in 2019.

### 283 3.4 VM Contextualization for particle physics applications

284 At the end of the boot process, the VMs are contextualized with the required software and the VM is configured  
 285 for the application software [17]. The contextualization is controlled by metadata that is passed to the “cloud-  
 286 init” process [18].

287 The contextualization of the VM instances for the ATLAS and Belle II experiments is very similar. Both  
 288 use the micro-CernVM virtual machine image [19], the CernVM File System (CVMFS) [20], and optionally,  
 289 *Shoal*, a dynamic web cache locator service [21].

290 The small micro-CernVM image (approximately 20 MB) saves storage space and can be started instantane-  
 291 ously. The image contains a read-only Linux kernel and a CernVM File System (CVMFS) client. An array  
 292 of squid proxy caches around the world are used to host frequently used files from CVMFS [22]. The CVMFS  
 293 client is configured to use a specific squid cache. We configure the CVMFS client to use the nearest squid cache  
 294 to the VM by querying Shoal.

295 The contextualization process then sets up the HTCondor client, and the environment for the experiment.  
 296 The accounting and monitoring processes are activated, and the CPU benchmark suite is run on the VM (see the  
 297 next section). Both the HTCondor client and the CPU benchmark are retrieved from CVMFS. The necessary  
 298 host certificates and SSH keys are added to the VM. All the VMs share the same host certificate; we currently  
 299 use certificates from GridCanada.

300 It is important to keep track of the resources delivered by a cloud to an experiment. We configure the VMs  
301 to write out their status every 15 minutes to Elasticsearch. The status includes configuration information, the  
302 fast benchmark and the timing information (obtained from `/proc/stat`). We are currently working on reporting  
303 our accounting numbers back into APEL for ATLAS [23]. The fast benchmark is discussed in the following  
304 section.

### 305 3.5 Operations and status

306 The CSV2 distributed compute cloud started production use in early 2019. The ATLAS production using the  
307 clouds in North America was shifted to CSV2, and tested extensively for many months. The Belle II production  
308 system was transitioned to CSV2 in June 2019. ATLAS and Belle II are now supported by a single instance of  
309 CSV2 but continue to use separate instances of HTCondor that are linked to their respective workload manage-  
310 ment systems.

311 In Fig. 5, we plot a histogram of the CPU hours weighted by a benchmark that is an estimate of the HEP-  
312 SPEC06 [24] benchmark. A centre usually runs the HEP-SPEC06 benchmark when their resources are being  
313 commissioned. As we are generally unaware of the underlying hardware, we can not use a static benchmark  
314 value. It is also impractical to run the HEP-SPEC06 on every VM, as it takes 2-3 hours (and there are also  
315 licensing issues). Instead, a fast-benchmark code (DB16) [25] is executed at the end of the contextualization of  
316 the VM instance to measure the performance of the dynamically provided resources.

317 Fig. 6 shows a histogram (normalized to unit area) of the fast-benchmark measured for every VM booted  
318 in 2019. The fast-benchmark gives a reproducible result if the VM is the only activity on the underlying node;  
319 however, the value is often underestimated if the node is busy or hyperthreading is used. The fast-benchmark  
320 is not ideal and there are ongoing efforts to find a more reliable estimate of the experiments' workload (see the  
321 presentation of the HEPiX Benchmark Working Group [26]).

322 In addition, we remark that the distributed compute cloud is using the Dynafed data federation service, de-  
323 veloped by CERN IT, for locating input data [27,28]. Dynafed federates existing storage systems and provides  
324 a link to the closest copy of the data on the basis of GeoIP information. Dynafed is used to federate existing  
325 storage of the ATLAS and Belle II experiments [29,30,31].

## 326 4 Summary

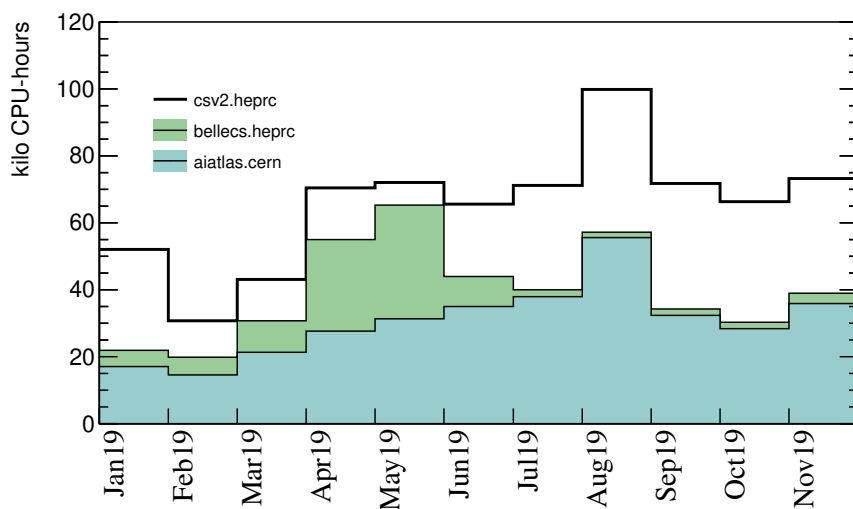
327 We have presented the design of a distributed cloud computing system based on the cloudscheduler VM pro-  
328 visioning platform. It remains a novel system for utilizing high-throughput workloads on multiple dedicated  
329 and opportunistic clouds located at remote locations and owned by other organizations. The system has pro-  
330 vided significant resources to the ATLAS and Belle II particle physics experiments. The cloudscheduler VM  
331 provisioning system has been updated to current software standards to make it more scalable and reliable, as  
332 well as adding new functionality. It is currently planned to use cloudscheduler for the computing part of the  
333 proposed Belle II Canadian Raw Data Centre. Other plans include further integration and use of the Dynafed  
334 data federator to expand the running of data-intensive applications.

335 The CSV2 code repository is hosted at GitHub (<https://github.com/hep-gc/cloudscheduler>) and documen-  
336 tation is found at [cloudscheduler.readthedocs.io](https://cloudscheduler.readthedocs.io).

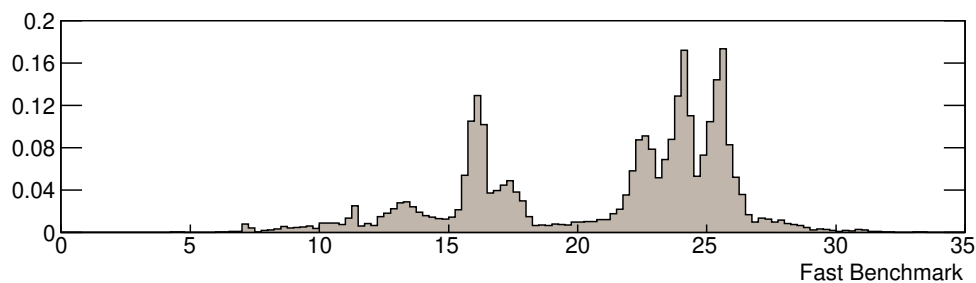
337 **Acknowledgements** We would like to thank the support of the Natural Sciences and Engineering Research Council of Canada  
338 and the Canadian Foundation of Innovation. In addition, the resources and generous support provided by Compute Canada, the  
339 Chameleon Project, the Edinburgh Compute and Data Facility (ECDF), Amazon and Microsoft are acknowledged.

## 340 Conflict of interest

341 The authors declare that they have no conflict of interest.



**Fig. 5** The fast-benchmark hours (x1000) per month in 2019. The (processor) hours are extracted from the kernel activity in `/proc/stat` and is typically 80-90% of the wall clock time. The blue histogram (aiatlas.cern) is the ATLAS HTCondor instance cloud hosted at CERN. The white histogram (csv2.heprc) and the green histogram (bellecs.heprc) are the ATLAS and Belle II HTCondor instances hosted in Victoria, Canada.



**Fig. 6** Histogram of the fast benchmark measured by the VMs. Note that the histogram is normalized to unit area. The fast-benchmark is an estimate of the HEP-SPEC06 benchmark. There are entries from multiple clouds. The types of processors in each cloud are unknown. When running on whole real CPUs the benchmark values show a narrow peak, while when using hyperthreaded CPUs the peaks are wider and shifted to lower values. The continuum entries between the peaks represent measured benchmark entries when other (unknown) processes are running on the same hypervisor.

342 **References**

- 343 1. E. Elsen, The end of computing's steam age,  
344 <https://home.cern/news/opinion/computing/end-computings-steam-age>
- 345 2. P. Armstrong *et. al.*, Cloud Scheduler: a resource manager for distributed compute clouds, arXiv preprint arXiv:1007.0050 (2010)
- 346 3. K. Keahey, M. Tsugawa, A. Matsunaga, and J. Fortes, Sky computing, *Internet Computing, IEEE*, 13(5):43–51, 2009,  
347 doi:10.1109/MIC.2009.94
- 348 4. R.P. Taylor *et. al.*, Consolidation of cloud computing in ATLAS, *J. Phys.: Conf. Ser.* 898 052008 (2017), doi:10.1088/1742-  
349 6596/898/5/052008
- 350 5. R.J. Sobie *et. al.*, Utilizing clouds for Belle II, *J. Phys.: Conf. Ser.* 664 022037 (2015), doi:10.1088/1742-6596/664/2/022037
- 351 6. A. McNab *et. al.*, Managing virtual machines with Vac and Vcycle, *J. Phys.: Conf. Ser.* 664 022031 (2015), doi:10.1088/1742-  
352 6596/664/2/022031
- 353 7. A. Tsaregorodtsev *et. al.*, DIRAC Distributed Computing Services, *J. Phys.: Conf. Ser.* 513 032096 (2014), doi:10.1088/1742-  
354 6596/513/3/032096
- 355 8. R. Grzymkowski *et. al.*, Belle II public and private cloud management in VMDIRAC, *J. Phys.: Conf. Ser.* 664 022021 (2015),  
356 doi:10.1088/1742-6596/664/2/022021
- 357 9. A. Amoroso *et. al.*, A modular (almost) automatic set-up for elastic multi-tenants cloud (micro)infrastructures, *J. Phys.: Conf.*  
358 *Ser.* 898 (2017) 082031, doi:10.1088/1742-6596/898/8/082031
- 359 10. B. Bockelman *et. al.*, Interfacing HTCondor-CE with OpenStack, *J. Phys.: Conf. Ser.* 898 092021 (2017), doi:10.1088/1742-  
360 6596/898/9/092021
- 361 11. K. Fransham *et. al.*, Research computing in a distributed cloud environment, *J. Phys.: Conf. Ser.* 256 (2010) 012003,  
362 doi:10.1088/1742-6596/256/1/012003
- 363 12. I. Gable *et. al.*, A batch system for HEP applications in a distributed IaaS cloud, *J. Phys.: Conf. Ser.* 331 062110 (2011),  
364 doi:10.1088/1742-6596/331/6/062110
- 365 13. D. Thain, T. Tannenbaum and M. Livny, Distributed computing in practice: the Condor experience, *Concurrency Computat.:*  
366 *Pract. Exper.*2005;17:323 (2005), doi:10.1002/cpe.938
- 367 14. The Worldwide LHC Computing Grid, <http://wlcg.web.cern.ch>
- 368 15. F.H. Barreiro Megino *et. al.*, PanDA for ATLAS distributed computing in the next decade *J. Phys.: Conf. Ser.* 898 052002  
369 (2017), doi:10.1088/1742-6596/898/5/052002
- 370 16. MariaDB open-source relational database, <https://mariadb.org>
- 371 17. K. Keahey and T. Freeman, Contextualization: Providing One-Click Virtual Clusters, *IEEE Xplore*: 06 January 2009, DOI:  
372 10.1109/eScience.2008.82
- 373 18. Cloud-init: the standard for customizing cloud instances. <https://cloud-init.io>
- 374 19. J. Blomer *et. al.*, Micro-CernVM: slashing the cost of building and deploying virtual machines, *J. Phys.: Conf. Ser.* 513 (2014)  
375 032009, doi:10.1088/1742-6596/513/3/032009
- 376 20. J. Blomer *et. al.*, New directions in the CernVM file system, *J. Phys.: Conf. Ser.* 898 062031 (2017), doi:10.1088/1742-  
377 6596/898/6/062031
- 378 21. I. Gable *et. al.*, Dynamic web cache publishing for IaaS clouds using Shoal, *J. Phys.: Conf. Ser.* 513 032035 (2014),  
379 doi:10.1088/1742-6596/513/3/032035
- 380 22. Squid is a caching proxy for the Web. <http://www.squid-cache.org>.
- 381 23. APEL is an accounting tool that collects accounting data from sites participating in the EGI and WLCG infrastructures,  
382 <https://wiki.egi.eu/wiki/APEL>
- 383 24. M. Michelotto *et. al.*, A Comparison of HEP code with SPEC1 benchmarks on multi-core worker nodes, *J. Phys.: Conf. Ser.*  
384 219 (2010) 052009, doi:10.1088/1742-6596/219/5/052009
- 385 25. P. Charpentier, Benchmarking worker nodes using LHCb productions and comparing with HEPspec06, *J. Phys.: Conf. Ser.* 898  
386 082011, doi:10.1088/1742-6596/898/8/082011
- 387 26. A. Valassi *et. al.*, Benchmarking WLCG resources using HEP experiment workloads, to be published in the proceedings of the  
388 Computing in High Energy Physics Conference, Adelaide, 2019.
- 389 27. Dynafed - The Dynamic Federation project. <http://lcgdm.web.cern.ch/dynafed-dynamic-federation-project>
- 390 28. F. Furano, O. Keeble and L. Field, Dynamic federation of grid and cloud storage, *Phys. Part. Nuclei Lett.* (2016) 13: 629.  
391 <https://doi.org/10.1134/S1547477116050186>
- 392 29. F. Berghaus *et. al.*, Federating distributed storage for clouds in ATLAS, *J. Phys.: Conf. Ser.* 1085 032027 (2018).  
393 doi:10.1088/1742-6596/1085/3/032027
- 394 30. M. Ebert *et. al.*, Using a dynamic data federation for running Belle-II simulation applications in a distributed cloud environment,  
395 EPJ Web of Conferences 214, 04026 (2019). <https://doi.org/10.1051/epjconf/201921404026>
- 396 31. F. Berghaus *et. al.*, The Dynafed data federator as grid site storage element, to be published in the proceedings of the Computing  
397 in High Energy Physics Conference, Adelaide, 2019.