

1 Using Kubernetes as an ATLAS computing site

2 *Fernando Harald Barreiro Megino*^{1,*}, *Jeffrey Ryan Albert*², *Frank Berghaus*², *Kaushik De*¹,
3 *FaHui Lin*¹, *Danika MacDonell*², *Tadashi Maeno*³, *Ricardo Brito Da Rocha*⁴, *Rolf Seuster*²,
4 *Ryan Paul Taylor*², and *Ming-Jyuan Yang*⁵ on behalf of the ATLAS collaboration

5 ¹University of Texas at Arlington, United States of America

6 ²University of Victoria, Canada

7 ³Brookhaven National Laboratory, United States of America

8 ⁴European Organization for Nuclear Research, Switzerland

9 ⁵Academia Sinica, Taiwan

10 **Abstract.** In recent years containerization has revolutionized cloud
11 environments, providing a secure, lightweight, standardized way to
12 package and execute software. Solutions such as Kubernetes enable
13 orchestration of containers in a cluster, including for the purpose of job
14 scheduling. Kubernetes is becoming a de facto standard, available at all
15 major cloud computing providers, and is gaining increased attention from
16 some WLCG sites. In particular, CERN IT has integrated Kubernetes into
17 their cloud infrastructure by providing an interface to instantly create
18 Kubernetes clusters, and the University of Victoria is pursuing an
19 infrastructure-as-code approach to deploying Kubernetes as a flexible and
20 resilient platform for running services and delivering resources.
21 ATLAS has partnered with CERN IT and the University of Victoria to
22 explore and demonstrate the feasibility of running an ATLAS computing
23 site directly on Kubernetes, replacing all grid computing services. We have
24 interfaced ATLAS' workload submission engine PanDA with Kubernetes,
25 to directly submit and monitor the status of containerized jobs. We
26 describe the integration and deployment details, and focus on the lessons
27 learned from running a wide variety of ATLAS production payloads on
28 Kubernetes using clusters of several thousand cores at CERN and the Tier
29 2 computing site in Victoria.

30 1 Introduction

31 The ATLAS experiment [1] has a long history of exploiting and adapting to new
32 technologies in cloud computing [2][3][4][5]. Containerization and Kubernetes are a
33 continuation of this trend in a new paradigm. The idea to use Kubernetes as a batch system
34 was conceived during the ATLAS-Google Data Ocean project [6], in which ATLAS
35 transferred data to Google Cloud Storage and then processed it using Google Compute
36 Engine virtual machines, according to the traditional Infrastructure-as-a-Service paradigm.
37 However, containers are a lightweight and more performant alternative to virtual machines,
38 and Kubernetes provides a robust, feature-rich and fully declarative platform for

* Corresponding author: barreiro [at] uta [dot] edu

39 automation and orchestration of containerized applications on any infrastructure.
40 Furthermore, Kubernetes could be an open-source replacement for traditional WLCG [7]
41 batch services that is widely adopted across the industry, available on all major cloud
42 providers, and provides a common interface across cloud providers and sites. The software
43 stack required at traditional WLCG sites could be considerably simplified by removing (or
44 encapsulating in containers) HEP-specific middleware. Moreover, since ATLAS has
45 already migrated to a fully container-based workload [8], it is simpler and more natural to
46 provide container-native infrastructure at sites, which can further obviate other setup steps
47 and enables new initiatives such as analysis preservation.

48 We have explored linking PanDA [9][10], ATLAS' Workload Management System
49 (WMS), with Kubernetes through the Harvester [11] resource interface. This contribution
50 will describe our initial integration model, ideas for improvement, and experience gained
51 using Kubernetes clusters at CERN and the University of Victoria (UVic).

52 **2 Site perspective: Kubernetes installation**

53 There are many methods and tools for provisioning Kubernetes clusters; a suitable one must
54 be chosen based on the circumstances and infrastructure at the site. At CERN, computing
55 resources are managed and provided primarily via one large Openstack cloud, and the
56 deployment of Kubernetes clusters as a service is integrated into the cloud with Magnum
57 [12]. At UVic, where a variety of resources are available including bare metal nodes and
58 several institutional clouds, Kubespray [13] was selected as a portable and flexible solution
59 for deploying Kubernetes clusters based on the Infrastructure-as-Code paradigm. Following
60 this approach, a cluster can be deployed or fully rebuilt from scratch within about 15
61 minutes.

62 ATLAS computing jobs rely on CVMFS [14], a distributed read-only file system
63 widely used in HEP, to access the experiment software. Two approaches have been used to
64 integrate CVMFS into Kubernetes clusters: at the host level underneath Kubernetes, and at
65 the application layer inside the cluster. In the host-based approach, a standard CVMFS
66 client is installed on each node in the cluster, and pods can mount CVMFS repositories
67 using a hostPath [15] volume. To ensure security, we enable the Kubernetes
68 PodSecurityPolicy admission controller and enforce a policy that whitelists the path /cvmfs
69 for read-only access. This approach avoids containerizing the CVMFS client and is suitable
70 if CVMFS is deemed to be essential infrastructure for a dedicated, purpose-built
71 Kubernetes cluster for HEP applications, as at UVic. On the other hand, at CERN the csi-
72 cvmfs driver [16] is used to provide pods with access to CVMFS, leveraging the
73 Kubernetes Container Storage Interface. This method contains CVMFS within the
74 application layer of the cluster, and is suitable on shared or general-purpose clusters or
75 commercial clouds, where it is not desirable or possible to modify the nodes in the cluster
76 with installation of domain-specific software such as CVMFS.

77 CVMFS can also be used to optimize the distribution of container images. Using
78 standard container runtimes, the initial startup time of a typical container is dominated by
79 pulling the image, even though only about 6% of the image data is actually read [17].
80 Containers that provide complex scientific software stacks like ATLAS' are often
81 significantly larger (≥ 1 GB), exacerbating the inefficiency of this approach. At UVic, we
82 integrated the CVMFS Docker Graph Driver plugin [18][19] into the Kubernetes cluster, so
83 that Docker loads image data on demand instead of pulling the entire image before starting
84 a container. In addition to accelerating initial container startup time by a factor of 4, this
85 saves significant amounts of bandwidth and removes a bottleneck for rapid and dynamic
86 scalability of the cluster, and integrates well with the host-based deployment of CVMFS.

87 To keep the cluster secure, at UVic we configured Role Based Access Control to allow
88 only those API actions which are suitable and necessary to run jobs on the cluster, and
89 applied another PodSecurityPolicy to forbid privileged pods. To monitor the status and
90 resource usage of the cluster, we use Prometheus [20], as shown in Figure 1.



91
92
93 **Fig. 1.** Prometheus monitoring showing resource usage for one of the Kubernetes clusters

94 **3 Central perspective: integration with the Workload** 95 **Management System through Harvester**

96 Harvester is the new resource-facing service developed by the PanDA team to interface the
97 central workload management system with distributed computing resources. It is flexible
98 and extensible, capable of exploiting HPCs, grid sites and clouds. Using new plugins
99 developed by our ASGC colleagues, Harvester can now also make use of Kubernetes
100 clusters.

101 The integration between Harvester and Kubernetes is illustrated in Figure 2. Harvester
102 uses the Kubernetes *Job* resource type to control pods. A *Job* is an abstraction of a pod that
103 is used to run non-persistent workloads, and ensures that a pod terminates successfully.
104 This relieves Harvester from having to track the status of individual pods and resubmit ones
105 that fail if e.g. a node experiences a problem. Harvester defines the specifications (e.g.
106 memory, CPU) for each *Job*, and Kubernetes schedules pods on nodes where the required
107 resources are available. The Harvester credential manager plugin regularly places fresh
108 X509 proxies in the Kubernetes cluster as secrets [21], which Kubernetes then makes
109 available to the pods so that pilot jobs can authenticate to grid services. We use the standard
110 ADC CentOS 7 container image available on Dockerhub [22]. Currently the container
111 image is statically defined per queue, but we plan to improve this so that it can be
112 dynamically defined on a per-job basis. Lastly, the pod runs a startup script and launches
113 the pilot, and the pilot retrieves a job from PanDA and executes it.

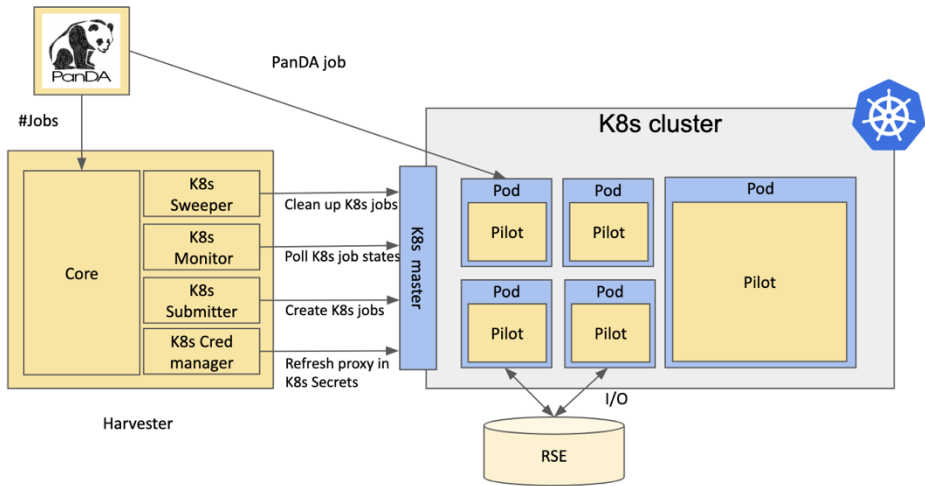


Fig. 2. Harvester-Kubernetes integration schematic.

4 Results

We have integrated Harvester with several Kubernetes clusters, most notably at CERN and UVic.

4.1 CERN

The first large-scale test was conducted at CERN in early 2019 on temporarily available resources. We created a 2000-core cluster consisting of 8-core nodes, and executed 1-core and 8-core ATLAS production jobs on the cluster. The main findings are depicted in Figure 3. Initially, we used the default Kubernetes scheduling algorithm, which balances load across nodes in a round-robin manner. This resulted in 1-core jobs being spread among nodes, blocking 8-core jobs from running on those nodes, thus decreasing the overall utilization of the cluster as more 1-core jobs run (see red ellipses in top image of Figure 3). To address this, we tuned the scheduling policy to pack the nodes, preferentially placing 1-core jobs on nodes that are already partially full rather than on empty nodes. This improved the resource usage of the cluster: there is only one small inefficiency when some nodes were draining 1-core jobs in order to start 8-core jobs (see grey ellipse in bottom image of Figure 3), but this is fundamentally unavoidable in any batch system. We are currently planning to repeat this exercise and potentially set up a permanent cluster at CERN.



136


137
138

139
140
141
142


Fig. 3. CPU usage on the CERN Kubernetes cluster with the default scheduling algorithm (top) vs the node-packing policy (bottom). Resources that are unused due to inefficient scheduling are indicated with red ellipses. Also note the cluster size decreased from 2000 cores to 1000 cores, due to issues with the csi-cvmfs driver and the infrastructure at CERN, which were addressed at the time.

143 **4.2 University of Victoria**

144 The second exercise was run on clusters at UVic. In order to evaluate and test several
 145 aspects of Kubernetes configuration, a small development cluster of 130 cores was
 146 deployed early on. The observations on this test cluster regarding Kubernetes scheduling
 147 were the same as at CERN (see Figure 4). The insights and solutions gained from the test
 148 cluster influenced the deployment of a larger, permanent T2 cluster with security and
 149 performance enhancements, which has been running stably without interventions since
 150 November 2019 (see Figure 5). 94% of wallclock time on this cluster is consumed by
 151 successful jobs, which will improve further once we address some known issues.
 152



153
154
155
Fig. 4. Default Kubernetes scheduling vs node packing scheduling on the UVic test cluster



156
157
158
159
Fig. 5. Stable operations of the 500-core UVic production cluster CA-VICTORIA-K8S-T2. The abrupt brief dips in usage are due to scheduled downtimes for network maintenance external to the site

160 **4 Conclusions and future work**

161 Kubernetes is the de facto standard for container orchestration, and many WLCG sites have
 162 expressed interest in using it. Kubernetes is usually used to manage applications and
 163 services, but can also be used directly as a batch system. We have demonstrated the use of
 164 Kubernetes, integrated with Harvester, as a native batch cluster for ATLAS production jobs
 165 at multiple sites, thereby significantly reducing the number of layers in the software and
 166 service stack of a traditional WLCG computing site.

167 The scale we have demonstrated so far is still smaller than ATLAS' typical grid
 168 usage, and some optimizations still need to be implemented. We are also interested in
 169 evaluating commercial cloud providers, and we have explored cluster federation, but have
 170 not yet found a satisfactory solution for this.

171 Furthermore, while this approach is very promising, we note that significant effort
 172 remains ahead to develop the maturity of the platform for our needs. The WLCG is a highly
 173 specialized infrastructure that has evolved over the last 15 years. Many implementation
 174 details have yet to be explored, such as accounting, traceability, fairshare-based scheduling,
 175 scaling the authentication and authorization model for wider use, and further adapting the
 176 workload and workflow to the container paradigm. More investigation, development and

177 evolution will be needed from the WLCG community to fully benefit from container-native
178 computing in the Kubernetes ecosystem.

179 Acknowledgements

180 This research was enabled in part by support provided by Compute Canada
181 (www.computecanada.ca), National Science Foundation (www.nsf.gov), US Department of
182 Energy (www.energy.gov) and WestGrid (westgrid.ca).

183 References

- 184 1. The ATLAS Collaboration, J. Inst. **3** S08003 (2008) doi:[10.1088/1748-](https://doi.org/10.1088/1748-)
185 [0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003)
- 186 2. F. Barreiro Megino et al., J. Phys. Conf. Ser. **396** 032011 (2012) doi:[10.1088/1742-](https://doi.org/10.1088/1742-)
187 [6596/396/3/032011](https://doi.org/10.1088/1742-6596/396/3/032011)
- 188 3. S. Panitkin et al., J. Phys. Conf. Ser. **513** 062037 (2014) doi:[10.1088/1742-](https://doi.org/10.1088/1742-)
189 [6596/513/6/062037](https://doi.org/10.1088/1742-6596/513/6/062037)
- 190 4. R. P. Taylor et al., J. Phys. Conf. Ser. **664** 022038 (2015) doi:[10.1088/1742-](https://doi.org/10.1088/1742-)
191 [6596/664/2/022038](https://doi.org/10.1088/1742-6596/664/2/022038)
- 192 5. R. P. Taylor et al., J. Phys. Conf. Ser. **898** 052008 (2017) doi:[10.1088/1742-](https://doi.org/10.1088/1742-)
193 [6596/898/5/052008](https://doi.org/10.1088/1742-6596/898/5/052008)
- 194 6. M. Barisits et al. on behalf of the ATLAS Collaboration, EPJ Web Conf. **214** 04020
195 (2019) doi:[10.1051/epjconf/201921404020](https://doi.org/10.1051/epjconf/201921404020)
- 196 7. LHC Computing Grid: Technical Design Report, document LCG-TDR-001, CERN-
197 LHCC-2005-024 (The LCG TDR Editorial Board) (2005)
- 198 8. A. Forti et al., Containers usage on the ATLAS grid infrastructure, 23rd International
199 Conference on Computing in High Energy and Nuclear Physics 2018, Sofia, Bulgaria
- 200 9. P. Nilsson et al. on behalf of the ATLAS Collaboration, J. Phys. Conf. Ser. **513** 032071
201 (2014) doi:[10.1088/1742-6596/513/3/032071](https://doi.org/10.1088/1742-6596/513/3/032071)
- 202 10. T. Maeno et al., J. Phys. Conf. Ser. **898** 052002 (2017)
- 203 11. T. Maeno et al., EPJ Web Conf., **214** (2019) 03030
- 204 12. OpenStack Magnum <https://docs.openstack.org/magnum/latest/>
- 205 13. Kubespray <https://kubespray.io/>
- 206 14. J. Blomer et al., J. Phys.: Conf. Ser. **898** 062031 (2017)
- 207 15. Kubernetes hostPath volumes
208 <https://kubernetes.io/docs/concepts/storage/volumes/#hostpath>
- 209 16. Kubernetes csi-cvmfs driver
210 <https://clouddocs.web.cern.ch/containers/tutorials/cvmfs.html>
- 211 17. T. Harter, B. Salmon, R. Liu, A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, 14th
212 USENIX Conference on File and Storage Technologies 194430 (2016)
- 213 18. N. Hardi, J. Blomer, G. Ganis, R. Popescu, J. Phys. Conf. Ser. **1085** 032019 (2018)
214 doi:[10.1088/1742-6596/1085/3/032019](https://doi.org/10.1088/1742-6596/1085/3/032019)
- 215 19. S. Mosciatti (2018) *Efficient container distribution at global scale*. Polytechnic
216 University of Milan, Milan, Italy. <https://hdl.handle.net/10589/144807>
- 217 20. Prometheus <https://prometheus.io/>

- 218 21. Kubernetes Secrets <https://kubernetes.io/docs/concepts/configuration/secret/>
- 219 22. ATLAS Distributed Computing CentOS7 image
- 220 <https://hub.docker.com/r/atlasadc/atlas-grid-centos7>