

# 1 Managing the ATLAS Grid through Harvester

2 *Fernando Harald Barreiro Megino*<sup>1,\*</sup>, *Aleksandr Alekseev*<sup>2</sup>, *Frank Berghaus*<sup>3</sup>, *David*  
3 *Cameron*<sup>4</sup>, *Kaushik De*<sup>1</sup>, *Andrej Filipcic*<sup>5</sup>, *Ivan Glushkov*<sup>1</sup>, *FaHui Lin*<sup>1</sup>, *Tadashi Maeno*<sup>6</sup>,  
4 and *Nicolo Magini*<sup>1</sup> on behalf of the ATLAS Collaboration

5 <sup>1</sup>University of Texas at Arlington, United States of America

6 <sup>2</sup>Tomsk Polytechnic University, Russia

7 <sup>3</sup>University of Victoria, Canada

8 <sup>4</sup>University of Oslo, Norway

9 <sup>5</sup>Jozef Stefan Institute, Slovenia

10 <sup>6</sup>Brookhaven National Laboratory, United States of America

11 <sup>7</sup>Iowa State University, United States of America

12 **Abstract.** ATLAS Computing Management has identified the migration  
13 of all computing resources to Harvester, PanDA's new workload  
14 submission engine, as a critical milestone for Run 3 and 4. This  
15 contribution will focus on the Grid migration to Harvester. We have built a  
16 redundant architecture based on CERN IT's common offerings (e.g.  
17 Openstack Virtual Machines and Database on Demand) to run the  
18 necessary Harvester and HTCondor services, capable of sustaining the load  
19 of O(1M) workers on the grid per day. We have reviewed the ATLAS Grid  
20 region by region and moved as much possible away from blind worker  
21 submission, where multiple queues (e.g. single core, multi core, high  
22 memory) compete for resources on a site. Instead we have migrated  
23 towards more intelligent models that use information and priorities from  
24 the central PanDA workload management system and stream the right  
25 number of workers of each category to a unified queue while keeping late  
26 binding to the jobs. We will also describe our enhanced monitoring and  
27 analytics framework. Worker and job information is synchronized with  
28 minimal delays to a CERN IT provided ElasticSearch repository, where we  
29 can interact with dashboards to follow submission progress, discover site  
30 issues (e.g. broken Compute Elements) or spot empty workers. The result  
31 is a much more efficient usage of the Grid resources with smart, built-in  
32 monitoring of resources.

33

34

---

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

## 35 **1 Introduction**

36 The Worldwide LHC Computing Grid (WLCG) [1] is a highly heterogeneous federation of  
37 computing sites with different middleware and increasingly special resources, such as  
38 Cloud or High Performance Computing (HPC) resources. PanDA [2] is the Workload  
39 Management System for ATLAS [3], managing all production and user jobs across the  
40 WLCG centers associated with the experiment. In order to exploit resources, PanDA is  
41 based on the Pilot paradigm [4], where Pilot jobs are submitted to the batch systems at sites.  
42 The Pilots retrieve the real payload from PanDA server and execute it.

43 Over the years numerous Pilot submission systems have been developed, frequently  
44 specializing on a certain subset of resources and having independent code bases. The  
45 Harvester project was born as an attempt to provide a universal Pilot submission system. At  
46 the same time, in the case of Grid resources, some improvements were needed to increase  
47 the stability and usage efficiency through a tighter integration with the PanDA Workload  
48 Management System allowing a more informed decision taking.

49 This contribution will focus on core Harvester design decisions and other significant  
50 aspects like new submission modes and monitoring. We will also show the process and  
51 results of migrating all Grid resources to Harvester.

## 52 **2 Harvester design decisions**

### 53 **2.1 Lightweight vs High Performance execution modes**

54 In order to be a universal service for any type of resource, Harvester needs to provide  
55 certain flexibility in terms of resource consumption and dependencies. In the case of HPCs,  
56 Harvester frequently needs to run on edge nodes with strict restrictions on installation and  
57 memory/CPU footprint of the application. In this case Harvester provides a lightweight  
58 mode, using a SQLite database and limiting the number of internally managed threads.

59 On the contrary for the Grid, central Harvester instances manage very large, costly  
60 infrastructures that need to be kept fully utilized. There are no important installation  
61 restrictions and it is preferable to host high performance services. In this case, Harvester  
62 typically uses a MySQL/MariaDB database and multiple Harvester processes can be started  
63 through frameworks like uWSGI [5].

### 64 **2.2 Fast integration of new resources**

65 To reduce the lead time until a new resource is exploited successfully, the Harvester code  
66 follows a plug-in approach: the code is split into a common core and a set of specific plug-  
67 ins that can be configured for each resource. Typically, new resources only require the  
68 implementation of resource-specific libraries for submission, monitoring and cleaning up  
69 workers that have finished. For most of the cases, these are python modules with lengths of  
70 a few hundred lines of code.

71 Up to date, plug-ins have been developed for HTCondor [6], ARC CE [7], Google  
72 Compute Engine [8], Kubernetes [9], SAGA [10], PBS [11], Cobalt [12] and Slurm [13].

73 Data management plugins can be also developed to handle input and output files. This  
74 is often required for HPCs, where the worker nodes have no external connectivity. However  
75 this is generally not required for the Grid, since the Pilot handles the data management.  
76  
77

78 **2.3 Queue unification**

79 ATLAS submits Pilots specifying the number of cores and memory for the jobs. In  
80 the past, sites provided multiple separate queues per job type:

- 81 • Analysis: usually single core payloads running with a non-production proxy security  
82 role
- 83 • Production: several different payloads summarized in Table 1.

84 **Table 1.** Requirements for the various production workloads

	<b>CPUs</b>	<b>Memory</b>
Single core	1	2 GB
Multi core	Depends on site, usually 8	2 GB/core
Single core, high memory	1	Depends on site, >2 GB
Multi core, high memory	Depends on site, usually 8	Depends on site, >2 GB/core

85  
86 An average site for ATLAS would have an Analysis queue, a single core and a multi-core  
87 production queue. Larger sites would also provide high memory queues. Pilots would be  
88 submitted to these queues independently, causing random competition for the slots and  
89 making it impossible to establish any control over the ratios.

90 In order to follow the Global Shares [14] priorities of ATLAS, it is more desirable to  
91 unify all queues at a site into one single queue that can accept Pilots with different sizes.  
92 These unified queues can be managed using either the Pull mode or the new Pull UPS mode  
93 (see next subsection).

94 **2.4 Submission modes supported by Harvester**

95 Harvester supports the classical push and pull workflows, and has extended the pull  
96 workflow for unified queues as follows:

- 97 • **Push (early binding):** the worker submitted to the batch queue is already assigned to a  
98 particular job and requests the CPU, memory, and potentially other requirements, of  
99 the specific job. The push workflow works natively with unified queues. The  
100 disadvantage is the early binding: once the worker is submitted, you can't control the  
101 queue time. If the queue time is long, a high priority job is stuck in the queue. Also,  
102 during the queueing time, a more important job might have appeared and needs to wait  
103 until the queue of previously submitted jobs has cleared.
- 104 • **Pull (late binding):** To circumvent the drawbacks of the push workflow, the pull  
105 workflow submits workers without any pre-assigned job. The advantage is that the  
106 WMS can choose the most important job right at the time the worker starts running.  
107 However, all workers submitted to the queue are the same and thus this mode does not  
108 support natively on unified queues. A dumb pilot submitter and multiple queues with  
109 different requirements on the same site causes competition between them and ATLAS  
110 can't control the ratio of workers across the queues.
- 111 • **Pull UPS (Unified Pilot Streaming):** This mode is an extension of the traditional pull  
112 mode to support unified queues. All types of jobs are queued together in the same  
113 queue and it's the WMS system deciding the fraction of workers of each type to  
114 submit, depending on the global priorities. The WMS decision is published to  
115 Harvester as a command.

116  
117

118

## 119 **3 Harvester monitoring**

### 120 **3.1 Worker monitoring**

121 Harvester reports the worker information up to the WMS, in our case PanDA. PanDA  
122 stores the worker information from the different Harvesters in one combined central Oracle  
123 table and also mirrors it to an ElasticSearch repository managed by CERN IT [15]. Our  
124 monitoring of choice follows the general trend to build dashboards on top of ElasticSearch,  
125 because of its speed and flexibility. We have built dashboards in Kibana for expert users  
126 and more user-friendly dashboards in Grafana for widespread use by site admins and  
127 shifters.

128 The dashboards allow the user to interact with the Harvester data at a detailed worker  
129 level and at high level overviews:

- 130 • The dashboards show a table with detailed worker information, which provide links to  
131 the logs and to the job submission file, to provide easy debugging of worker failures.
- 132 • The dashboards also provide several plots to see the submission evolution broken down  
133 by different criteria (Harvester instance, resource type, state, etc.). There are also plots  
134 to see the worker distribution across the Computing Elements of a site.

### 135 **3.2 Service monitoring**

136 Our service monitoring collects important information about the Harvester instances:

- 137 • Disk, CPU and memory usage
- 138 • Worker submission, update and completion rates

139 These metrics allow to easily identify when an instance is in trouble or misbehaving.  
140 We have an alerting system in place that sends a mail to the central Harvester team in case  
141 a metric has exceeded a threshold. The alerting system has proven very useful and warns us  
142 quickly about issues. These alerting systems are critical in times of low operational  
143 manpower and shifters in ATLAS.

### 144 **3.2 Site monitoring**

145 We are also implementing monitoring to identify broken sites. There are views in place  
146 grouping broken sites and broken Computing Elements, based on high ratios of workers in  
147 bad states (failed or cancelled submission). We show also the error messages happening on  
148 those sites.

149 We are working on a monitoring to identify inactive sites, where the submission rate  
150 is lower than what would be expected. There can be different reasons for this, including  
151 misbehavior of Harvester submission rate calculation for the site. In the future we would  
152 like to automate actions, like issuing a reset of the worker submission calculation or  
153 submitting tickets to the site.

## 154 **4 Migration of the Grid to Harvester**

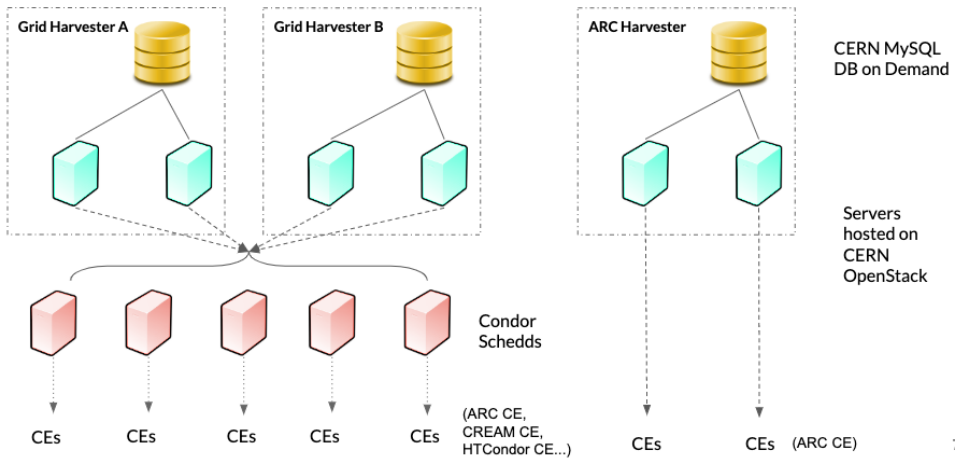
155 **4.1 Central infrastructure**

156 We have setup a central high-performance Harvester infrastructure for the whole Grid  
157 migration. It is based on CERN IT provided services. The databases are hosted by the DB  
158 on Demand [16] project, while the servers are virtual machines from the OpenStack service.  
159 We have distributed the Grid sites across three Harvester instances:

- 161 • Harvester A: US, CERN Tier 0, France, Russia, Canada, Netherlands
- 162 • Harvester B: CERN, Germany, UK, Taiwan, Italy, Spain
- 163 • Harvester ACTA: ARC CE sites

164 Each instance consists of two active machines and one shared database for each  
165 instance. The two machines provide redundancy and are in different zones of the CERN  
166 computing center. If one machine is unavailable, the other machine can take over the full  
167 load.

168 Non-ARC computing elements rely on HTCondor as interface. There are five load-  
169 balanced HTCondor machines for the whole grid.




170

171 **Fig. 1.** Harvester infrastructure overview

172 **4.2 Migration process**

173 The whole grid was migrated June 2018 to March 2019. During the migration process we  
174 watched stability and scaling of the services, improved the efficiency of the HTCondor  
175 interfaces, simplified the configuration of queues through the ATLAS Grid Information  
176 System (AGIS) [2017] and improved some interactions with the database.



177

178  
179

**Fig. 2.** Overview of the Grid migration progress to Harvester, in number of workers. This histogram is extracted from our Harvester worker monitoring.

180  
181  
182  
183  
184

During the first period we migrated production queues to Harvester. Although production jobs occupy the better part of the grid slots, they are rather easy to handle from a Harvester load point of view: they run for multiple hours and typically occupy multiple cores. During the migration a major effort was devoted to unify production queues to the pull UPS mode explained in section 2.4.

185  
186  
187  
188

Once the infrastructure was proven to be scalable and stable, we also migrated analysis queues to Harvester. Analysis jobs represent less than 20% of the grid slots, but they are short and run on a single core. From a Harvester load point of view, they are heavier and require a much higher worker submission rate.

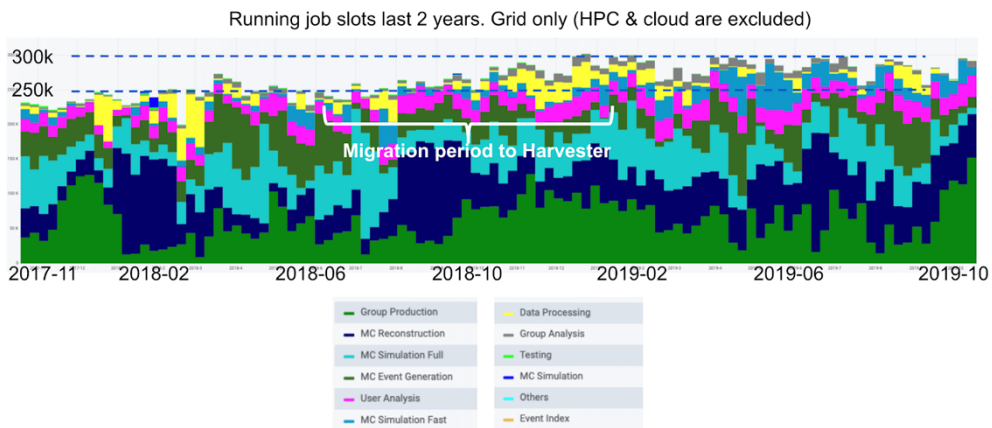
## 189 5 Results

190  
191  
192  
193  
194  
195

Harvester is a project that took 2 years from initial discussion to full roll out. It is contributing to a better usage of the ATLAS Grid. During the roll out of Harvester, we observed a significant usage of job slots. While there are multiple potential contributions happening in parallel (e.g. 3-4% of pledge increase or increased usage of the ATLAS Tier 0), we believe that Harvester has made a significant contribution to the increase for two reasons:

196  
197  
198  
199  
200

- Harvester is more aggressive than previous Pilot factories in worker submission and competes better on shared sites across multiple experiments.
- Through the queue unification and the pull UPS submission, the worker submission is also more intelligent. It avoids uninformed competition between ATLAS queues at the same site and follows better the ATLAS job priorities.



201

202 **Fig. 3.** Running job slots on Grid resources before and after the migration to Harvester

203 Unified queues for production are working very well and we are currently working on  
 204 further unifying production and analysis queues, so that a standard Grid site only needs to  
 205 provide one queue for all job types.

206 Lastly, we want to enhance automation of issues on the Grid to reduce the need of  
 207 human operational effort.

## 208 Acknowledgements

209 This research was enabled in part by support provided by National Science Foundation  
 210 ([www.nsf.gov](http://www.nsf.gov)) and US Department of Energy ([www.energy.gov/](http://www.energy.gov/)).

## 211 References

- 212 1. LHC Computing Grid: Technical Design Report, document LCG-TDR-001, CERN-  
 213 LHCC-2005-024 (The LCG TDR Editorial Board) (2005)
- 214 2. T. Maeno et al. *J. Phys. Conf. Ser.* **898** 052002 (2017)
- 215 3. ATLAS Collaboration 2008 The ATLAS Experiment at the CERN Large Hadron  
 216 Collider *J. Inst.* **3** S08003
- 217 4. P. Nilsson et al. *J. Phys. Conf. Ser.* **513** 032071 (2014)
- 218 5. uWSGI <https://uwsgi-docs.readthedocs.io/en/latest/>
- 219 6. HTCondor <https://research.cs.wisc.edu/htcondor/>
- 220 7. ARC CE <http://www.nordugrid.org/arc/ce/>
- 221 8. Google Compute Engine <https://cloud.google.com/compute/>
- 222 9. Kubernetes <https://kubernetes.io/>
- 223 10. SAGA <https://saga-python.readthedocs.io/en/latest/>
- 224 11. PBS <https://www.nas.nasa.gov/hecc/support/kb/121/>
- 225 12. Cobalt <https://trac.mcs.anl.gov/projects/cobalt>
- 226 13. Slurm <https://www.slurm.schedmd.com/>
- 227 14. F. Barreiro Megino et al. *EPJ Web Conf.*, **214** (2019) 03025

- 228 15. P. Saiz et al. EPJ Web Conf., (to be published)
- 229 16. R. Aparicio et al. J. Phys. Conf. Ser. **664** 052021 (2015)
- 230 17. A. Anisenkov et al. J. Phys. Conf. Ser. **664** 062001 (2015)