

1 Production Experience with the ATLAS Event 2 Service

3 **D Benjamin¹, P Calafiura², T Childers⁶, K De³, W Guan⁴,**
4 **T Maeno⁵, P Nilsson⁵, V Tsulaia², P Van Gemmeren⁶ and**
5 **T Wenaus⁵ on behalf of the ATLAS Collaboration**

6 ¹Duke University, 134 Chapel Drive, Durham, NC 27708, USA

7 ²Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA

8 ³University of Texas at Arlington, 701 South Nedderman Drive, Arlington, TX 76019, USA

9 ⁴University of Wisconsin, 1150 University Avenue, Madison, WI 53706, USA

10 ⁵Brookhaven National Laboratory, PO Box 5000, Upton, NY 11973, USA

11 ⁶Argonne National Laboratory, 9700 S. Cass Ave, Argonne, IL 60439, USA

12 E-mail: VTsulaia@lbl.gov

13 **Abstract.** The ATLAS Event Service (AES) has been designed and implemented for efficient
14 running of ATLAS production workflows on a variety of computing platforms, ranging from
15 conventional Grid sites to opportunistic, often short-lived resources, such as spot market
16 commercial clouds, supercomputers and volunteer computing. The Event Service architecture
17 allows real time delivery of fine grained workloads to running payload applications which process
18 dispatched events or event ranges and immediately stream the outputs to highly scalable Object
19 Stores. Thanks to its agile and flexible architecture the AES is currently being used by
20 grid sites for assigning low priority workloads to otherwise idle computing resources; similarly
21 harvesting HPC resources in an efficient back-fill mode; and massively scaling out to the 50-100k
22 concurrent core level on the Amazon spot market to efficiently utilize those transient resources
23 for peak production needs. Platform ports in development include ATLAS@Home (BOINC) and
24 the Google Compute Engine, and a growing number of HPC platforms.

25 After briefly reviewing the concept and the architecture of the Event Service, we will
26 report the status and experience gained in AES commissioning and production operations on
27 supercomputers, and our plans for extending ES application beyond Geant4 simulation to other
28 workflows, such as reconstruction and data analysis.

29 1. Introduction

30 The ATLAS Experiment [1] processes its data at about 140 computing centers around the world
31 at a scale of about 4M CPU-hours/day. To date it has accumulated a globally distributed data
32 volume in excess of 220 Petabytes. Even with such a massive processing scale, the experiment
33 is resource limited. The ATLAS physics program can benefit from applying more compute
34 resources to Monte Carlo simulation, and over the next decade the situation will become even
35 more critical because the LHC [2] and ATLAS upgrade programs will bring an order of magnitude
36 increase in computing requirements. In view of the steady demand for new computing resources,
37 it becomes very important for the experiment to not only efficiently use all CPU power available
38 to it, but also to proactively leverage opportunistic computing resources.



39 Opportunistic computing resources have a large potential for expanding the ATLAS
40 processing pool. Such resources include cost-effective clouds such as the Amazon spot market [3],
41 supercomputers (HPCs), shared Grid resources and volunteer computing (ATLAS@Home) [4].
42 Porting of regular ATLAS workloads (e.g. simulation, reconstruction) to opportunistic resources
43 does not come for free. In order to use them fully and efficiently ATLAS has implemented
44 a fine-grained event processing system - the ATLAS Event Service (AES) [5] - in which the job
45 granularity changes from input files to individual events or event ranges. The Event Service
46 delivers fine-grained workload to the running event processing application (the payload) in real
47 time. After processing each event range (about 10 min processing time), the Event Service writes
48 the corresponding output into a separate file and saves the output file into a secure location, such
49 that Event Service jobs can be terminated practically at any time with minimal data losses. This
50 architecture allows the Event Service to efficiently adapt to the characteristics of opportunistic
51 resources, in which a job slot lifetime is unpredictable and may be either very short or very long.

52 In order to efficiently utilize CPU resources of supercomputers, we have developed an
53 HPC-specific implementation of the Event Service called Yoda [6], which leverages MPI for
54 running massively parallel event processing jobs on multiple HPC compute nodes simultaneously.
55 Yoda has been developed and prepared for production usage on the Edison supercomputer at
56 the National Energy Research Scientific Computing Center (NERSC), Berkeley, USA. Since late
57 2015 Yoda has been running ATLAS simulation production workloads at NERSC and in 2016
58 it delivered about 20M CPU hours to the experiment.

59 In section 2 of this paper we describe the concept and the architecture of the Event Service.
60 Yoda is described in Section 3 and the AES commissioning status is presented in Section 4.
61 During the commissioning phase of Yoda we studied various factors which can have a visible
62 effect on the CPU efficiency of compute nodes. Such factors include initialization time of
63 the payload application, sequential running of several payloads on a compute node within
64 the same MPI-submission, and handling of fine-grained outputs. The former two factors are
65 discussed in Section 5, while in Section 6 we present the results of our studies of the performance
66 of Object Stores, which are used by the Event Service as an intermediate storage for fine-grained
67 outputs produced by payload applications.

68 **2. The ATLAS Event Service**

69 The JEDI [7] (Job Execution and Definition Interface) extension to PanDA [8] adds new
70 functionality to the PanDA server to dynamically break down tasks in a way that optimally
71 utilizes available processing resources. With this capability, tasks can be broken down at
72 the level of either individual events or event clusters (ranges). This functionality allowed us
73 to develop the ATLAS Event Service capable to dynamically deliver to a compute node only
74 that portion of the input data which will be actually processed there by the payload application.
75 Input data is streamed to the compute node in real time in small portions. While the payload
76 persists, it can elastically continue to consume new inputs and stream away outputs with no
77 need to tailor workload execution time to resource lifetime. A schematic view of the Event
78 Service workflow is shown in Figure 1.

79 On the compute node the PanDA Pilot establishes a connection with the PanDA server over
80 HTTP and starts a parallel event processing application (payload) in order to utilize all available
81 CPU cores. The payload application in the Event Service is represented by AthenaMP [9],
82 a process-parallel version of the ATLAS data processing framework Athena. AthenaMP starts
83 as a serial process, which first goes through the application initialization phase, then forks several
84 event processors (workers) and informs the pilot that it is ready for data processing. The pilot
85 downloads event range identifiers (strings) from the PanDA server and delivers them in real
86 time to the running AthenaMP application, which assigns them to its workers on a first-come,
87 first-served basis. The worker uses the event range string to locate the corresponding input file

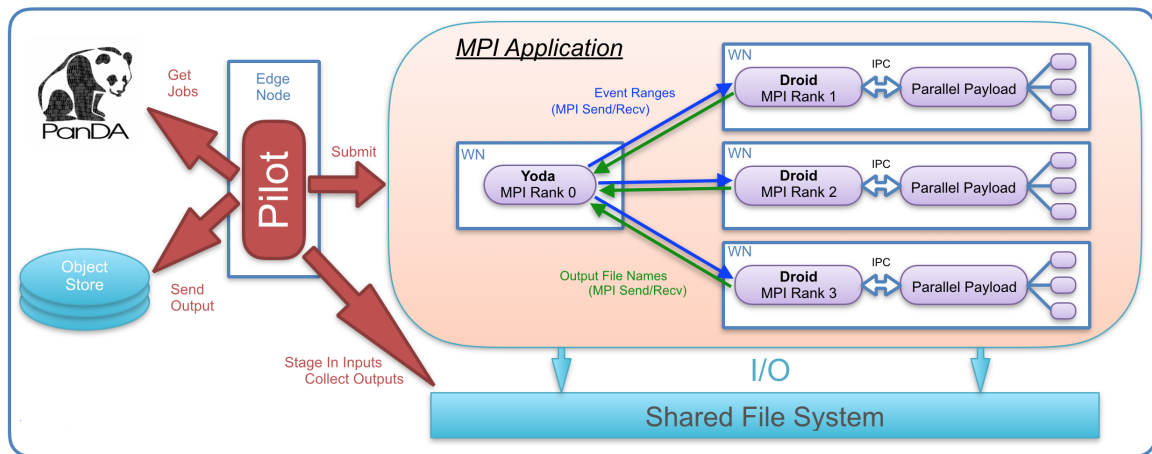


Figure 2. Schematic view of Yoda

115 definitions from the PanDA server and streams out the outputs produced by Yoda jobs to
 116 the Object Store. Yoda applications implement the master-slave architecture in which rank 0
 117 is the master and all other ranks are the slaves. For the development of Yoda ranks we reused
 118 the code of the conventional Event Service and implemented lightweight versions of PanDA
 119 JEDI (hence Yoda, a diminutive Jedi) and the PanDA Pilot (Droid). Yoda (rank 0) orchestrates
 120 the entire MPI-application by continuously distributing fine-grained workloads to Droids (rank
 121 N , $N \neq 0$) and collecting their outputs. On a compute node Droid starts the payload application
 122 and delivers the workload to it exactly the same way the PanDA Pilot does it on compute nodes
 123 of the conventional Event Service applications. This allows us to run the same configuration of
 124 AthenaMP payload on HPC and on other Event Service platforms such as the Grid and Clouds.
 125 The outputs produced by AthenaMP on the compute nodes are temporarily stored to the HPC
 126 Shared File System until the Pilot streams them out to Object Stores.

127 4. Event Service commissioning

128 We have chosen ATLAS Geant4 Simulation [10, 11] as a first use-case for the Event Service in
 129 general and for Yoda in particular. Simulation jobs use a substantial fraction of the ATLAS CPU
 130 budget on the Grid which makes it very beneficial for the experiment to offload its simulation
 131 to other computing platforms such as opportunistic resources and HPCs. On the other hand,
 132 simulation jobs are CPU-intensive with minimal I/O requirements and relatively simple handling
 133 of in-file metadata, characteristics which allowed us to make rapid progress in the development
 134 of the Event Service and Yoda components and to begin commissioning the Event Service for
 135 production usage.

136 Until recently NERSC supercomputers (Edison and Cori Phase I) had been our primary
 137 platforms for the development and commissioning of the AES. We started to run Simulation
 138 production workloads with the Event Service (Yoda) on the Edison HPC in late 2015, and in
 139 2016 Yoda delivered about 20M CPU-hours to the ATLAS collaboration. Also in late 2015
 140 we successfully scaled Event Service up to 50,000 concurrent processors on the Amazon Spot
 141 Market cloud. In Summer 2016 the Event Service commissioning effort was shifted over to Grid
 142 sites and it has been showing steady progress since then. Event Service deployment on volunteer
 143 computing (ATLAS@Home) has not progressed significantly due to manpower shortages.

144 5. Performance studies

145 During the commissioning of Yoda on the Edison supercomputer we studied various factors
146 which can have a visible effect on the CPU efficiency of Yoda ranks. In this section we discuss
147 payload initialization time and sequential running of several payloads on a compute node within
148 the same MPI submission.

149 5.1. Payload initialization

150 During its initialization step AthenaMP reads a large number of files from the disk. These files
151 include python scripts, shared libraries, XML configuration files, static replicas of the geometry
152 and conditions database, etc. If the ATLAS offline software release is installed on the HPC
153 shared file system, then the concurrent reading of software installation files by many compute
154 nodes during the payload initialization phase can lead to a serious performance bottleneck.
155 For example, we have observed rather poor scaling of AthenaMP initialization time on Edison
156 compute nodes when all instances of AthenaMP were accessing a software release installed on
157 Edison's scratch file system (Lustre).

158 In order to work around this problem we package the entire ATLAS software release into
159 a single tarball. At the beginning of its execution the Droid first unpacks this tarball into
160 the memory-resident disk on the compute node, then it starts AthenaMP and lets it initialize
161 on the local copy of the software release. With this approach we eliminate concurrent reading of
162 the shared release installation by all Yoda payloads which considerably speeds up
163 the initialization phase of the entire Yoda application.

164 Although with this mechanism we achieved very good scaling up to 1,000 concurrent starts,
165 the preparation of software release tarballs requires considerable manual effort and so is not
166 considered sustainable in the long run. On the Cori Phase I supercomputer we studied
167 AthenaMP initialization performance scaling by installing software releases on different systems
168 including Lustre, Burst Buffer [12] and Shifter [13]. So far the results obtained with the Shifter
169 system look the most promising.

170 5.2. Sequential running of multiple payloads on the same compute node

171 Before submitting Yoda jobs to the HPC batch system, the Pilot first needs to get the workload
172 from a PanDA production task. This mechanism is illustrated by Figure 3. PanDA tasks
173 consist of many jobs and each job requires processing of many events. When a new task gets
174 defined in PanDA all its jobs contain the same number of events. Depending on the number
175 of ranks (compute nodes) allocated for a given Yoda job, the Pilot decides how many PanDA
176 jobs should be processed by this MPI-job and passes this information over to Yoda. Yoda then
177 assigns each PanDA job to one or more ranks. The strategy here is to keep each compute node
178 busy for the entire lifetime of the MPI-job. In cases when Yoda does not have enough time to
179 process all events from a PanDA job, all leftover events are returned back to the PanDA server,
180 which generates new PanDA jobs containing only these leftover events. This mechanism leads to
181 the creation of many PanDA jobs with a number of events less than the task's default number
182 per job.

183 If Yoda has to process PanDA jobs with a small number of events, it assigns several such
184 jobs to a single compute node. The Droid running on this compute node deals with multiple
185 PanDA jobs in sequence, which means several instances of AthenaMP are started and stopped
186 by the Droid during its lifetime. While AthenaMP is going through the initialization phase,
187 all CPU cores on the node are idling and in this way significant CPU time is wasted. Dealing
188 with multiple PanDA jobs within a single PanDA task in this way quite often leads to rather
189 poor overall CPU efficiency of Yoda jobs. In the future we plan to overcome this problem
190 by implementing a new concept of Jumbo Jobs in PanDA. With Jumbo Jobs each production
191 task in PanDA will be represented by a single PanDA job. Thus, Yoda will not have to deal

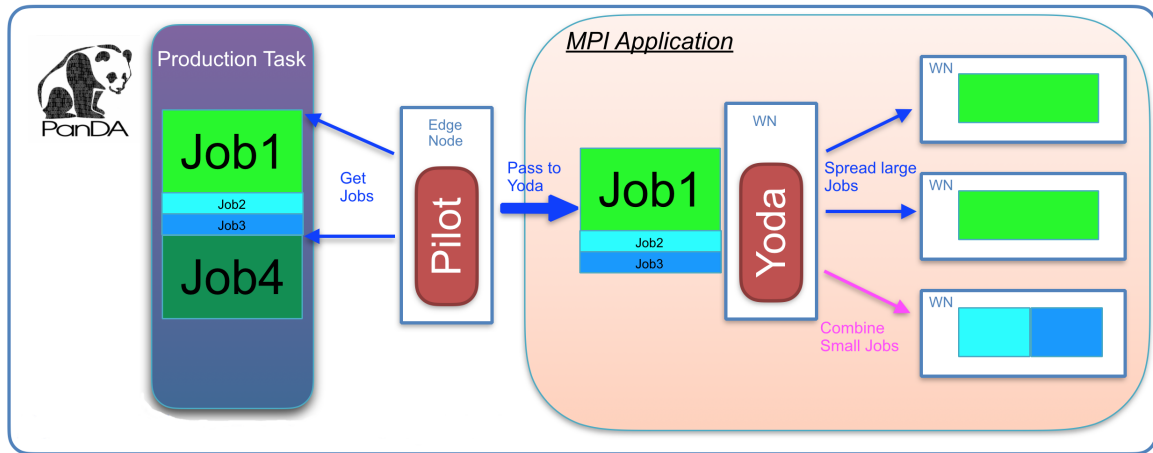


Figure 3. Yoda dealing with multiple PanDA jobs




Figure 4. CPU efficiency of Yoda compute nodes

192 with multiple PanDA jobs and no time will be wasted initializing more than one instance of
 193 AthenaMP on a single compute node.

194 5.3. CPU efficiency of Yoda compute nodes

195 Figure 4 shows three time-line plots demonstrating CPU efficiency of Yoda compute nodes.
 196 These plots were obtained from Yoda test runs on Edison. The X axis of each plot shows
 197 the wall time in minutes since the beginning of Droid execution on the compute node, and
 198 each bin on the Y axis corresponds to one CPU core (Edison compute nodes have 24 physical
 199 CPU cores). The white color on the plot means the core is idle, turquoise means the core is
 200 processing an event and red means event processing was started but not finished for some reason
 201 (e.g. segmentation fault occurred, or the job was killed because it reached its wall time limit).

- 202 • The plot on the left is an example of poor CPU efficiency caused by the very long
 203 initialization time of AthenaMP;
- 204 • The plot in the middle is an example of poor CPU efficiency caused by running more than
 205 one PanDA job on a single compute node;
- 206 • The plot on the right is an example of good CPU efficiency: just one PanDA job runs on
 207 the compute node, initialization is fast and the number of events is enough to keep the node
 208 busy for the entire job lifetime.




Figure 5. OS bandwidth dependency on the object size

209 6. Interaction with Object Stores

210 Intermediate output files produced by the Event Service payload applications are shipped in real
 211 time to the Object Stores (OS). PanDA then generates specialized jobs which merge these files
 212 into final outputs. Such merge jobs usually run on Grid sites. The initial implementation of
 213 Yoda was sending event range outputs directly from Edison compute nodes to the Object Store
 214 at BNL. The stage-out process was coupled with the event processing, therefore data transfer
 215 issues (e.g. network connection problems, slow file upload) were affecting the CPU efficiency
 216 of Yoda compute nodes. To avoid these problems we decoupled data transfer to the OS from
 217 event processing on the compute nodes, making output uploading the responsibility of the Pilot,
 218 which runs on the HPC edge node.

219 As part of the Event Service commissioning at NERSC we studied the Object Store
 220 performance by running a series of tests which involved uploading of objects of different size to
 221 the CEPH OS at BNL. We observed that the clients can overload the OS with various errors
 222 occurring including authentication errors, inability to connect to a bucket, inability to write an
 223 object, long running writer, etc. This suggests that either the client software should have retry
 224 and perhaps queuing capabilities, or we need a server side system that can regulate OS writes.

225 Another important observation is that we can achieve much higher bandwidth by increasing
 226 the object sizes. This is demonstrated on Figure 5, which shows that by grouping 10-50 events
 227 into a single transfer (transfer size 6.8-34 MB) we achieved 950MB/s upload speed vs 125MB/s
 228 for single event transfers (transfer size 0.68 MB).

229 7. Summary

230 The Event Service has been commissioned to run ATLAS Geant4 Simulation production on
 231 HPC systems. The commissioning process on Grid sites is well underway and other deployment
 232 platforms (e.g. clouds, volunteer computing) are expected to follow.

233 Several important lessons were learned during the development and testing of Yoda at
 234 NERSC:

- 235 (i) Primary causes of sub-optimal usage of CPU resources on the compute nodes are slow
 236 initialization of the payload and the fact that for the time being Yoda must combine multiple
 237 PanDA jobs into a single MPI-submission;
- 238 (ii) By staging out large numbers of small files we can saturate Object Stores;

239 (iii) Data stage out must be decoupled from event processing.

240 By addressing the issues listed above we were able to successfully scale production Yoda jobs
241 up to 700 compute nodes (almost 17,000 cores) on Edison HPC at NERSC.

242 In the future we plan to further develop Event Service functionality by implementing
243 the Event Streaming Service. Also we will be applying the Event Service to other ATLAS
244 production workflows beyond Geant4 Simulation (e.g. Reconstruction and Analysis) with
245 the ultimate goal to make the Event Service a unified workflow architecture across all ATLAS
246 computing platforms.

247 8. Acknowledgments

248 The results presented in this paper have been obtained by using resources of the National Energy
249 Research Scientific Computing Center, a DOE Office of Science User Facility supported by
250 the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-
251 05CH11231.

252 References

- 253 [1] ATLAS Collaboration, 2008 *JINST* **3** S08003
- 254 [2] L. Evans and P. Bryant LHC Machine, 2008 *JINST* **3** S08001
- 255 [3] The Amazon Elastic Computing Cloud, <http://aws.amazon.com/ec2/>
- 256 [4] Adam-Bourdarios C et al. on behalf of the ATLAS Collaboration 2015 ATLAS@Home: Harnessing Volunteer
257 Computing for HEP *J. Phys.: Conf. Ser.* **664** 022009
- 258 [5] Calafiura P et al. on behalf of the ATLAS Collaboration 2015 The ATLAS Event Service: A new approach
259 to event processing *J. Phys.: Conf. Ser.* **664** 062065
- 260 [6] Calafiura P et al. on behalf of the ATLAS Collaboration 2015 Fine grained event processing on HPCs with
261 the ATLAS Yoda system *J. Phys.: Conf. Ser.* **664** 092025
- 262 [7] De K, Golubkov D, Klimentov A, Potekhin M and Vaniachine A on behalf of the ATLAS Collaboration 2014
263 Task Management in the New ATLAS Production System *J. Phys.: Conf. Series* **513** 032078
- 264 [8] Maeno T for the ATLAS Collaboration 2008 PanDA: Distributed production and distributed analysis system
265 for ATLAS *J. Phys.: Conf. Series* **119** 062036
- 266 [9] Calafiura P et al. on behalf of the ATLAS Collaboration 2015 Running ATLAS workloads within massively
267 parallel distributed applications using Athena Multi-Process framework (AthenaMP) *J. Phys.: Conf. Ser.*
268 **664** 072050
- 269 [10] GEANT4 Collaboration, S. Agostinelli et al., 2003 *Nucl. Instrum. Meth. A* **506** 250
- 270 [11] ATLAS Collaboration 2010 ATLAS Simulation Infrastructure *Eur. Phys. J* **C70** 823
- 271 [12] Bhimji W et al. 2016 Extreme I/O on HPC for HEP using the Burst Buffer at NERSC. Proceedings of the
272 CHEP2016 conference *J. Phys.: Conf. Ser.*
- 273 [13] Gerhardt L et al. 2016 Using Shifter to Bring Containerized CVMFS to HPC. Proceedings of the CHEP2016
274 conference *J. Phys.: Conf. Ser.*