

# Large Scale Software Building with CMake in ATLAS

J Elmsheuser<sup>1</sup>, A Krasznahorkay<sup>2</sup>, E Obreshkov<sup>3</sup>, A Undrus<sup>1</sup>  
on behalf of the ATLAS Collaboration

<sup>1</sup> Brookhaven National Laboratory, USA

<sup>2</sup> CERN, Switzerland

<sup>3</sup> University of Texas, Arlington, USA

E-mail: [Attila.Krasznahorkay@cern.ch](mailto:Attila.Krasznahorkay@cern.ch)

**Abstract.** The offline software of the ATLAS experiment at the Large Hadron Collider (LHC) serves as the platform for detector data reconstruction, simulation and analysis. It is also used in the detector's trigger system to select LHC collision events during data taking. The ATLAS offline software consists of several million lines of C++ and Python code organized in a modular design of more than 2000 specialized packages. Because of different workflows, many stable numbered releases are in parallel production use. To accommodate specific workflow requests, software patches with modified libraries are distributed on top of existing software releases on a daily basis. The different ATLAS software applications also require a flexible build system that strongly supports unit and integration tests. Within the last year this build system was migrated to CMake.

A CMake configuration has been developed that allows one to easily set up and build the above mentioned software packages. This also makes it possible to develop and test new and modified packages on top of existing releases. The system also allows one to detect and execute partial rebuilds of the release based on single package changes. The build system makes use of CPack for building RPM packages out of the software releases, and CTest for running unit and integration tests.

We report on the migration and integration of the ATLAS software to CMake and show working examples of this large scale project in production.

## 1. Introduction

The ATLAS experiment's [1] offline software is responsible for reconstructing, simulating and analysing petabytes of data collected by the detector at the Large Hadron Collider. It is composed of multiple "projects" that are each built up out of many "packages".

As the software is built out of millions of lines of C++ and Python code, developers cannot build the entire thing when developing software in the collaboration. They are instead just building a select number of packages against numbered releases, or nightly builds of the full software release, overriding just the libraries/executables/etc. provided by the package from the ones found in the "base release".

The software is mainly distributed to the users through CVMFS [2], but it is also important to be able to install the software on an offline computing node with as little effort as possible.

ATLAS, together with LHCb, used the CMT [3] build system for building its offline software until recently, but the experiment has been looking for a build system with much wider community support to replace CMT since a few years. Finally at the beginning of 2016 the



40 build of the software release meant to reconstruct and simulate data in 2017 and beyond has  
41 been switched over to be built by CMake [4].

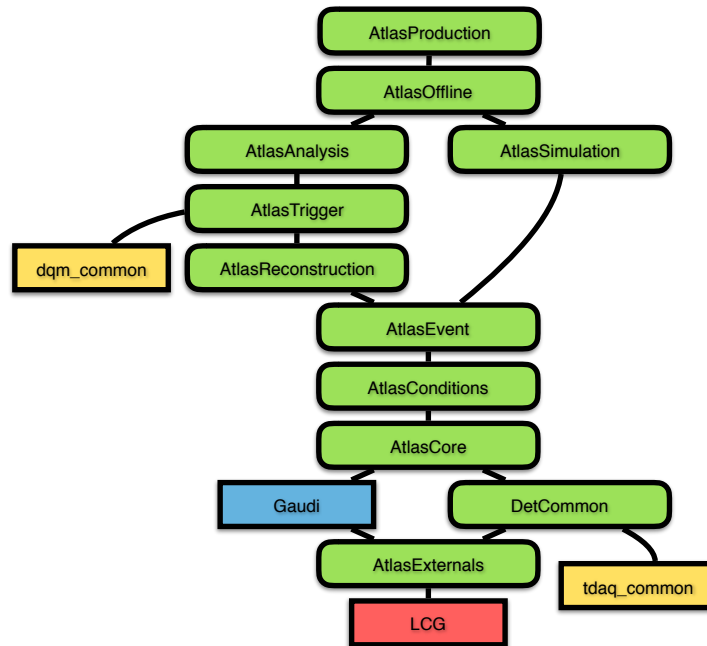
## 42 2. Software organisation

43 ATLAS’s software is broken up into  $\mathcal{O}(2000)$  “packages”, which can provide:

- 44 • Shared libraries to be used by other packages;
- 45 • Module libraries used by the software framework to load software components;
- 46 • Executables performing specialised tasks;
- 47 • Scripts, configuration files, etc. to be installed alongside the build targets of the package.

48 A package declares explicitly which other packages it requires for its own build, what software  
49 products it needs to build, and what already existing files it needs to install.

50 A software “project” is just a collection of packages, built in an order consistent with the  
51 dependency declarations in the package configurations. The projects themselves depend on each  
52 other similarly to how packages do. The packages in a certain project are allowed to depend only  
53 upon packages that are in a project that their project depends on. To demonstrate the highest  
54 project complexity used, the project configuration of the software used for reconstructing events  
55 in ATLAS is shown in Figure 1.



**Figure 1.** Projects and their relationships in the ATLAS offline software. “LCG” represents the externals provided by the LCG software release [5], the yellow boxes represent ATLAS online software built with their own configuration, and Gaudi [6] is built using its own CMake configuration.

## 56 3. The implementation

57 While the concept of packages is taken from the CMT build system, we wanted to keep the same  
58 structure with CMake as well. We used CMake’s subdirectory concept to keep the configuration  
59 of the CMT packages separated with CMake. In this setup we have a master CMakeLists.txt  
60 file describing the project, which needs to:

- 61 • Set up some basic properties of the project, like its name and version;
- 62 • Find the base project(s) that it depends on;
- 63 • Find all subdirectories in the source tree that have a `CMakeLists.txt` file in them, and
- 64 include them as subdirectories of the software project;
- 65 • Set up all variables necessary for using CTest and CPack with the project.

66 All of the common code used by the project and package configuration code is collected into  
67 two special packages in the source tree:

- 68 • AtlasCMake provides a large number of functions and macros all sharing the `atlas_` prefix,
- 69 which help in performing common tasks in the project and package configuration files;
- 70 • AtlasLCG collects the code necessary for finding all non-ATLAS-specific externals for the
- 71 project. It provides `FindFoo.cmake` CMake modules for finding these externals, and setting
- 72 up dependencies on the RPM packages of these externals for the RPM package built from
- 73 the ATLAS project.

### 74 *3.1. Package configuration*

75 The configuration of an average package may look something like the following.

```

76 # The name of the package:
77 atlas_subdir( ExamplePackage )
78
79 # The packages that this package depends on:
80 atlas_depends_on_subdirs(
81     PUBLIC
82     Control/AthenaKernel
83     PRIVATE
84     Control/CxxUtils )
85
86 # External packages needed for the build:
87 find_package( ROOT COMPONENTS Core Hist )
88 find_package( Boost COMPONENTS regex )
89
90 # Build a shared library:
91 atlas_add_library( ExampleLibrary ExamplePackage/*.h src/*.cxx
92     PUBLIC_HEADERS ExamplePackage
93     INCLUDE_DIRS ${ROOT_INCLUDE_DIRS}
94     PRIVATE_INCLUDE_DIRS ${Boost_INCLUDE_DIRS}
95     LINK_LIBRARIES ${ROOT_LIBRARIES} AthenaKernel
96     PRIVATE_LINK_LIBRARIES ${Boost_LIBRARIES} CxxUtils )
97
98 # Build an application:
99 atlas_add_executable( ExampleApp util/ExampleApp.cxx
100     LINK_LIBRARIES ExampleLibrary )
101
102 # Add tests for the package:
103 atlas_add_test( ExampleLibrary_test
104     SOURCES test/ExampleLibrary_test.cxx
105     LINK_LIBRARIES ExampleLibrary )
106 atlas_add_test( ExampleApp_test
107     SCRIPT test/ExampleApp_test.sh )
108
109 # Install files from the package:
110 atlas_install_python_modules( python/*.py )

```

111 The `atlas_` prefixed functions take care of setting up the build and installation of components  
112 according to some common rules used in the build. They also take care of declaring the shared  
113 libraries to be exported, so that child projects can make use of them.

114 Every package declares two helper targets:

- 115 • `Package_PkgName` builds all library and executable targets in the package;
- 116 • `Package_PkgName_tests` builds all of the test executables declared in the package.

117 They are meant to help during development, when the user set up the build of many packages  
118 at the same time.

### 119 *3.2. Project configuration*

120 As discussed earlier, the project's main configuration is done in its main `CMakeLists.txt` file,  
121 which is constructed in the following way:

```
122 # The minimum required CMake version:
123 cmake_minimum_required( VERSION 3.2 FATAL_ERROR )
124
125 # Find this project's main dependency. To pick up its CMake configuration:
126 find_package( AtlasEvent )
127
128 # Set up the flags for CTest:
129 atlas_ctest_setup()
130
131 # Set up the project, and all of the packages in it:
132 atlas_project( AtlasReconstruction 21.0.5
133     USE AtlasEvent 21.0.5
134     FORTRAN )
135
136 # Set up the flags for CPack:
137 atlas_cpack_setup()
```

138 By finding the base project, the ATLAS CMake helper code from the base project is included  
139 into the current project, making it capable of using the `atlas_` prefixed functions.

140 The `atlas_project` function is responsible for doing the heavy lifting in the project's  
141 configuration. It:

- 142 • Sets up all general build flags for the project, including the output directories for all file  
143 types inside of the build directory;
- 144 • Looks up all of the packages in the source directory tree, and includes them with CMake's  
145 built-in `add_subdirectory` call;
- 146 • After having included all packages, it includes the shared libraries from its base project(s)  
147 as imported targets.

148 *3.2.1. Usage of exported targets* The build system has to allow us to patch a release by building  
149 a package on top of the release that the release itself has also built. Jobs executed in this  
150 environment would pick up libraries/executables/installed files from the patched package.

151 In order to do this every project exports all of its shared libraries during the project's  
152 installation, adding the name of the project as a prefix to the exported library's name. When  
153 calling `find_package` on the project, all of these imported targets become visible to the  
154 configuration already. Package developers are strongly advised against using such imported  
155 targets directly, though, since a given package should not have to know what other packages it

156 is compiled together with in the same project. They should only refer to shared libraries that  
157 they need for the build without any project name prefixes.

158 To make the build work, the `atlas_project` function calls `find_package` on the base project  
159 after having processed all packages, in the following way:

```
160 find_package( AtlasEvent 21.0.5 COMPONENTS INCLUDE QUIET )
```

161 What the base project does in this case is that it loops over all of its imported targets, and  
162 checks if a target with the same name is already declared in the current project. If it is, the  
163 project assumes that its library is being re-built in the current project, and lets that target be  
164 used by every other component in the project. If a target with that name is not defined, then  
165 it makes a copy of that imported target, removing the project name prefix from its name.

### 166 *3.3. Runtime environment setup*

167 The ATLAS offline software uses a large number of software products that are not developed  
168 by ATLAS, and are not part of the base operating system that we use the offline software on.  
169 Some of these externals we build as part of the offline software, but most of them we pick up  
170 for the build from custom locations.

171 Most of the used externals are picked up from software bundles provided by the LCG project.  
172 CMake finds those externals for the build using the `AtlasLCG` code described earlier.

173 Unfortunately CMake does not provide a built-in way for setting up a custom environment for  
174 the built/installed project, as its base assumption is that the build results should work without  
175 any special setup on the build host. This is however not true in our case. We have to ensure  
176 on our own that libraries and executables that the offline software used during the build are  
177 available in the runtime environment of the software.

178 This is done by the build generating simple shell scripts for setting up the necessary runtime  
179 environment. At the end of the project configuration we iterate over all external packages  
180 that were found during the configuration of the project, and generate a setup script that saves  
181 information on how to extend the environment for running our software.

182 In order to simplify the environment setup when building multiple projects, these setup  
183 scripts know how to find the setup scripts of the current project's base project. The user only  
184 needs to use the script from the highest project that he/she wants to use, and the script takes  
185 care of executing the setup from all of the base projects of that project as well.

### 186 *3.4. Making the projects relocatable*

187 As the projects must be installable in any location, we need to be careful with setting up both  
188 the CMake configuration files generated during the CMake installation, and the environment  
189 setup scripts generated by our private code. Neither of them can hold any absolute path names  
190 in order to make the projects relocatable.

191 As we have full control over the generation of the environment setup scripts, there all that  
192 we do is to define a small number of environment variables that, when defined before using the  
193 setup script, direct the code to the correct directory to set up the externals from.

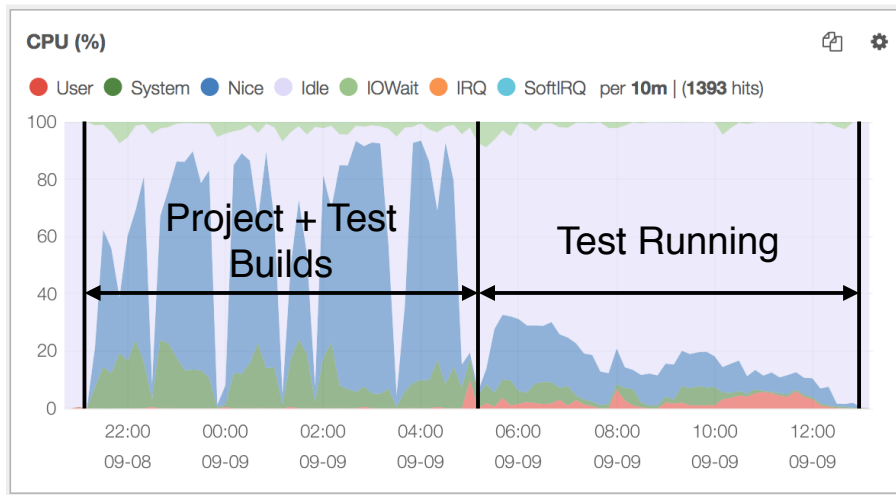
194 Making the CMake project files generated during the installation step is a bit more difficult.  
195 Since the file generated by CMake cannot be made relocatable out of the box, we use a custom  
196 script that processes the CMake generated code, and makes it relocatable. We make use of the  
197 following formalism in the CMake configuration:

```
198 install( SCRIPT ${CMAKE_BINARY_DIR}/atlas_export_sanitizer.cmake )
```

199 **4. Performance**

200 One of the goals of the build system migration was to speed up the build of the ATLAS offline  
201 software, and streamline the installation procedure of the software after it was built. CMake  
202 helps with this, as it can parallelize the build of independent components very efficiently.

203 The build machine monitoring system tells us that the build makes quite efficient use of the  
204 resources on the build machines. See Figure 2.



**Figure 2.** The CPU utilisation of the build machine while building the full ATLAS offline software.

205 In this project configuration, which we kept for a transition period between the build systems,  
206 the build times became just a little shorter. The single process execution steps necessary between  
207 the individual projects can be seen as clear downtimes during the build in Figure 2.

208 In order to simplify the offline software builds of ATLAS, we are now merging most of the  
209 projects into a single one called *Athena*, which project’s code will be stored in a single Git  
210 repository. In such a setup, on a fast build machine, it is possible to build the full (single)  
211 project in  $\mathcal{O}(3)$  hours, hitting the original performance goal of the build system migration.

212 **5. Conclusions**

213 In an effort to streamline its software build procedures, and use a system in common with many  
214 High Energy Physics Computing projects, ATLAS has switched to using CMake to describe  
215 the build procedures of its offline/simulation/trigger/analysis software. We are now in the final  
216 stages of validating the new, CMake built offline software of the experiment for the 2017 data  
217 taking.

218 The new build system, together with migrating the experiment’s software to Git, makes the  
219 development procedures followed by the ATLAS software developers much more in common  
220 with practices used in the software development community on the whole.

221 The new system has so far met the requirements of the experiment both in features provided,  
222 and performance. And so it is expected to be kept for LHC’s Run 3 and beyond.

223 **References**

224 [1] The ATLAS Collaboration 2008 *JINST* **3** S08003  
225 [2] Blomer J *et al* 2015 The evolution of global scale filesystems for scientific software distribution *Computing*  
226 *in Science and Engineering* **17(6)** 61-71  
227 [3] Arnault C 2001 Experiencing CMT in software production of large and complex projects *Proc. Int. Conf. on*  
228 *Computing in High Energy and Nuclear Physics* CHEP 2001 (Beijing, China)

- 229 [4] Martin K and Hoffman B 2007 An open source approach to developing software in a small organization *IEEE*  
230 *Software* **24** Number 1
- 231 [5] <http://cern.ch/lcgsoft>
- 232 [6] Clemencic M *et al* 2010 Recent developments in the LHCb software framework Gaudi *J. Phys.: Conf. Ser.*  
233 **219** 042006