



Contribution ID: 374

Type: Oral

Fast distributed compilation and testing of large C++ projects

Tuesday, November 5, 2019 4:45 PM (15 minutes)

High energy physics experiments traditionally have large software codebases primarily written in C++ and the LHCb physics software stack is no exception. Compiling from scratch can easily take 5 hours or more for the full stack even on an 8-core VM. In a development workflow, incremental builds are often not sufficient for quick compilation on a typical PC (e.g. due to changes to headers or rebasing), and large shared servers are not practical as users have no control and maintenance is an issue. Even though support for building partial checkouts on top of published project versions exists, by far the most practical development workflow involves full project checkouts because of off-the-shelf tool support (git, intellisense, etc.)

This paper details a deployment of distcc, a distributed compilation server, on opportunistic resources such as development machines. The best performance operation mode is achieved when preprocessing remotely and profiting from the shared CernVM File System. A 10 (30) fold speedup of elapsed (real) time is achieved when compiling Gaudi, the base of the LHCb stack, when comparing local compilation on a 4 core VM to remote compilation on 80 cores, where the bottleneck becomes non-distributed work such as linking. Compilation results are cached locally using ccache, allowing for even faster rebuilding. A recent distributed memcached-based shared cache is tested as well as a more modern distributed system by Mozilla, sccache, backed by S3 storage. These allow for global sharing of compilation work, which can speed up both central CI builds and local development builds. Finally, we explore remote caching and execution services based on Bazel, and how they apply to Gaudi-based software for distributing not only compilation but also linking and even testing.

Consider for promotion

No

Primary author: MATEV, Rosen (CERN)

Presenter: MATEV, Rosen (CERN)

Session Classification: Track 5 – Software Development

Track Classification: Track 5 – Software Development