GlideinBenchmark: collecting resource information to optimize provisioning

Marco Mambelli, Fermilab; Shrijan Swaminathan, Purdue University

GlideinBenchmark

GlideinBenchmark is a Web application leveraging the pilot infrastructure of GlideinWMS to benchmark resources, and collects and publishes data to automate the optimal selection of resources.
GlideinBenchmak has two components:

- the Runner: service controlling the GlideinWMS Factory to execute the desired benchmark on the desired resources
- the Viewer: Web dashboard serving via a RESTful interface and displaying the benchmark results

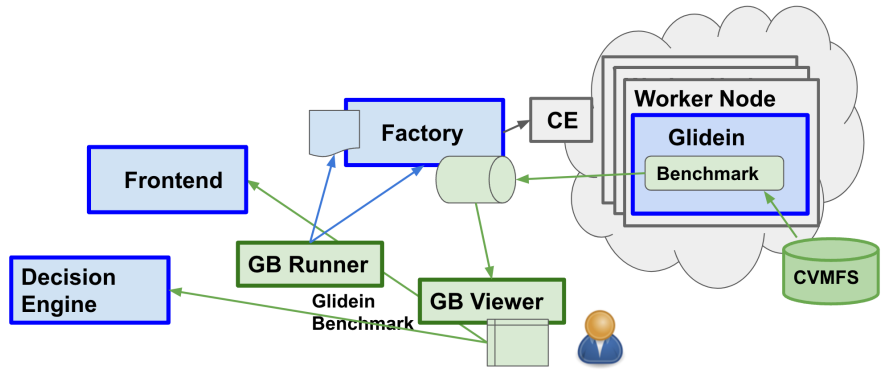

GlideinBenchmark architecture

Manual benchmarks

Manual benchmarks have been useful for provisioning resources, especially for instance selection in cloud resources, but are time-consuming to perform. Therefore the idea to automate the process via GlideinBanchmark

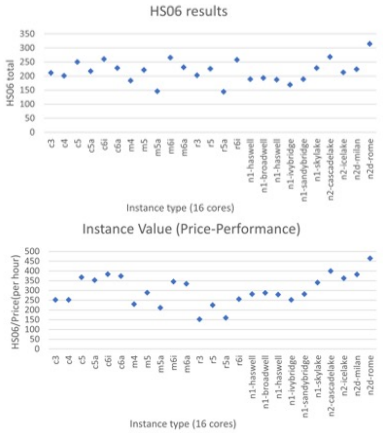

Benchmark results - Higher generation CPUs perform better, even more when the price-performance ratio is considered

Benchmark collection and integration

GlideinWMS provides the framework to run the benchmarks:

- The Glidein, pilot job, runs the benchmark in containers via Apptainer, an unprivileged container platform, and sends back the results
- The Factory runs automated benchmarks via the Glidein custom-script mechanism – used normally for tests and setups – and on-demand benchmarks as one-off Glidein submissions
- The Factory also collects from the Glidein the results, then stored and served by the GlideinBenchmark Viewer

Root-less benchmarks are selected and saved as container images on CVMFS, a distributed file system, to make them available to the Glideins on all resources

The HEPCloud Decision Engine can use the benchmark results to evaluate its Figure Of Merit, a numeric value associated to each resource and used to decide where to provision new computing resources.

The GlideinWMS Frontend will take advantage of the benchmark results as well when these are integrated in its provisioning heuristics used to request Glideins to the Factory.

This manuscript (FERMILAB-POSTER-24-0305-CSAID) has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.
This work was done in part under the DOE Omni Technology Alliance Internship Program by ORISE