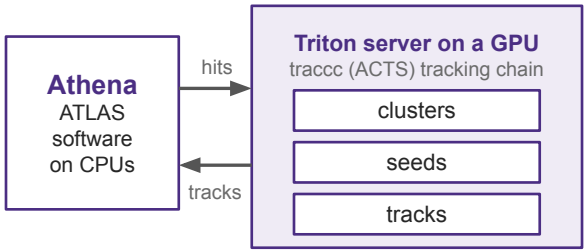
Keeping GPU Tracking-as-a-Service in Sync with Athena Through Shared Memory

Owen Huang[1], Yuan-Tang Chou[2], Miles Cochran-Branson[1], Xiangyang Ju[3], Beojan Stanislaus[3], Vakho Tsulaia[3], Yulei Zhang[1], Shih-Chieh Hsu[1]

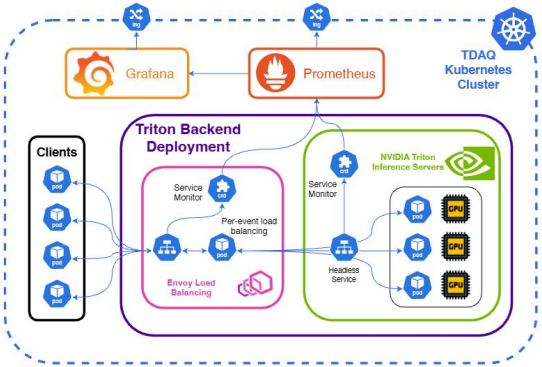
[1]University of Washington [2]National Tsing Hua University [3]Lawrence Berkeley National Laboratory

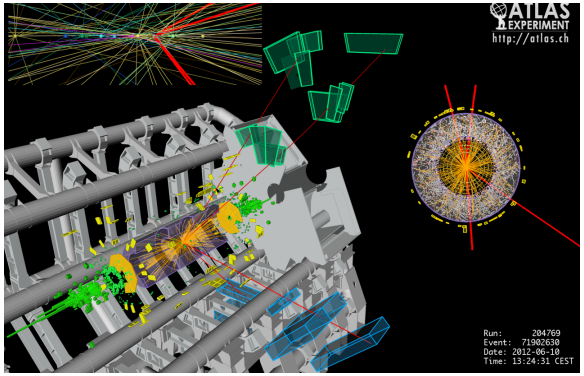
CPAD 2026, Seattle

1 Athena, Triton and traccc

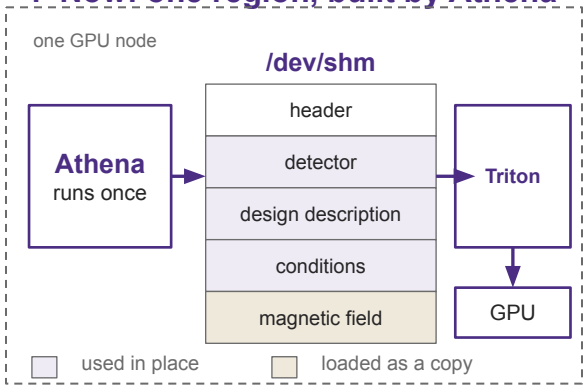


2 Tracking as a service

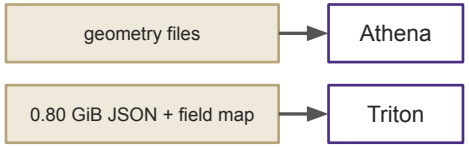




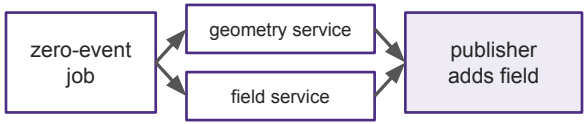
4 Now: one region, built by Athena



3 Before: two copies that can drift



5 One zero-event Athena job



6 Results

13.0 s
model load
was 29.1 s

246 MB
geometry in memory
was 0.80 GiB

7 Next

- Throughput benchmarking
- Physics validation