

Interactive parallel analysis in ALICE experiment at CERN: ALICE Analysis Facility

Arsen Hayrapetyan, Yerevan Physics Institute, Lab. #143, YerPhi/ALICE group
ahairape@mail.yerphi.am, Arsen.Hayrapetyan@cern.ch

Outline

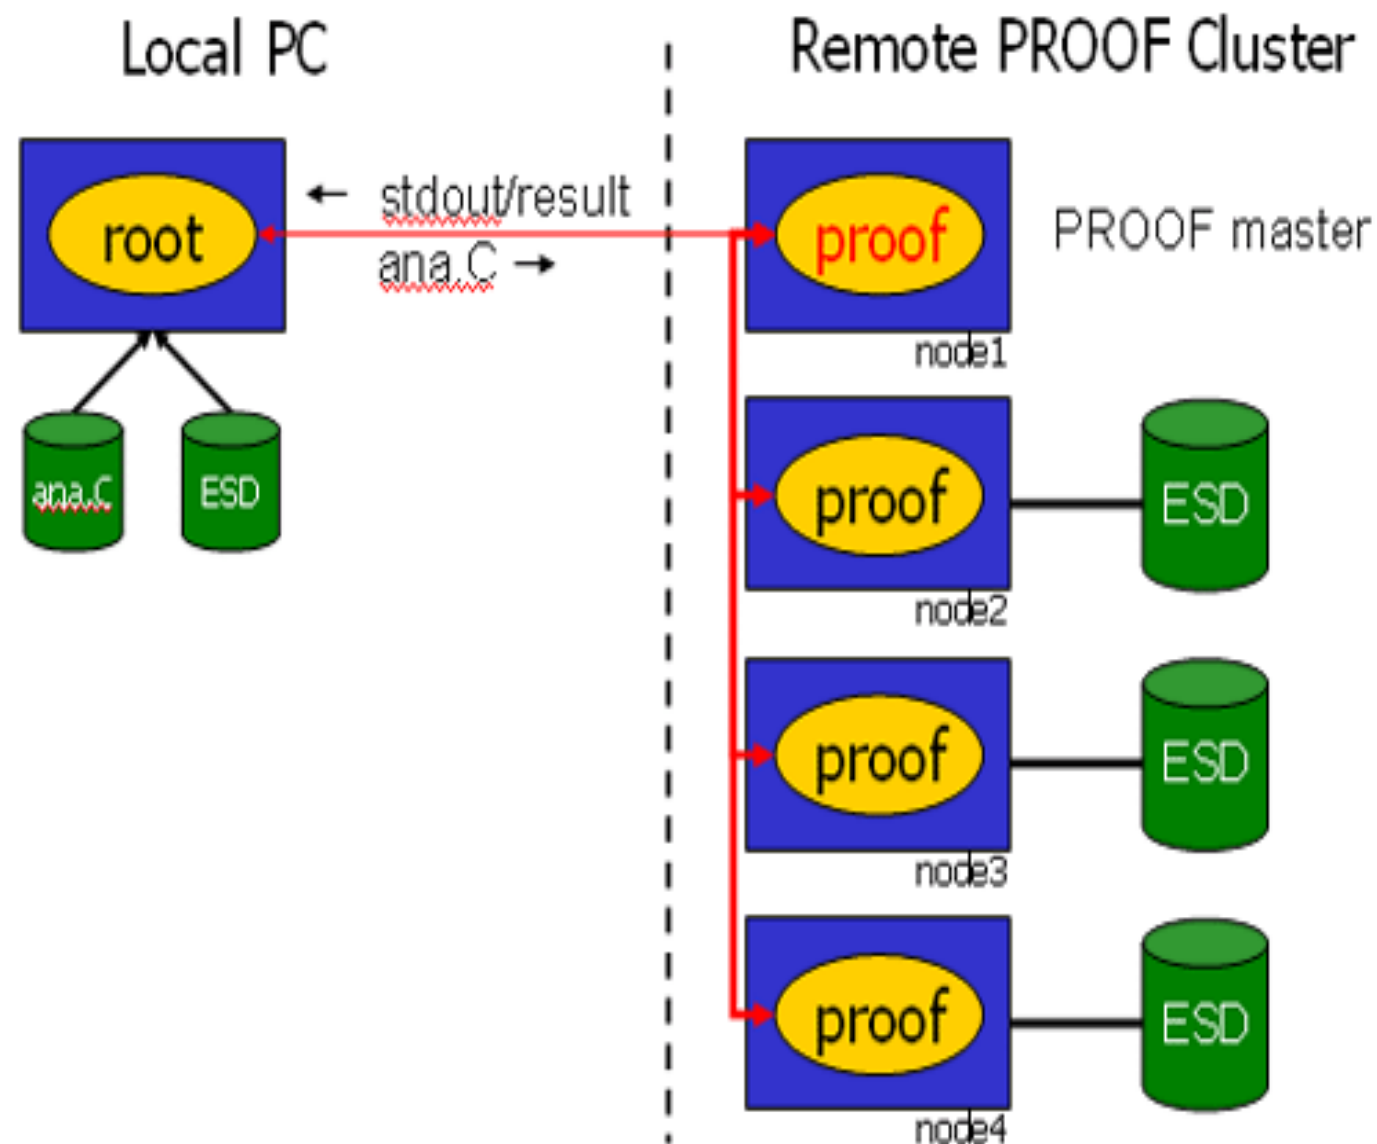
- ALICE Analysis Facility at CERN: CAF
- Tools to build the CAF
 - Parallel processing of events: PROOF
 - Data management: xrootd, afdsmgrd
 - Software packages distribution: alien (packman)
- What is CAF for?
 - Comparison to grid and local analysis
- Monitoring
- Multiple analysis facilities: AAF
- Statistics on CAF
- Summary
- Documentation and user support

ALICE Analysis Facility at CERN: CAF

- ALICE Analysis Facility at CERN is called CAF (CERN Analysis Facility)
 - Cluster of machines
 - 208 cores (26 x 8)
 - 80 TB of storage
 - Runs the following software:
 - **PROOF** (**P**arallel **ROOT** Facility) for *parallel processing of data chunks*
 - **xrootd** (**e**xtended **r**oot **d**aemon) and **afdsmgrd** (**A**nalysis facilities **d**ataset **m**anager **d**aemon) for data management (*copying, files from grid to CAF storage - staging, removing the files from CAF storage, etc.*)
 - **alien** (**A**LICE **e**nvironment) for *installing and updating experiment analysis software (AliRoot, ROOT, geant)*

PROOF computing model (simple setup, no submerging)

PROOF Schema



- User runs analysis task from her laptop/desktop and connects to cluster via “master node” (within root session; proof is part of root)
- Proof starts N=58 (current default setting for max. number of workers on CAF) workers on “slave” nodes
- The data set to be analysed is divided into N subsets
- Each worker processes one dataset and produces its output
- The output is sent to master node and merged
- Merged output is sent to the user

Data management: afdsmgrd, xrootd

- The physics data to be analysed is normally copied to CAF storage from grid beforehand. The process of copying is called staging.
 - It is possible to access those data directly on grid during analysis, but this slows down the analysis.
- To copy the data afdsmgrd is used which, in turn, uses xrootd to perform the actual copy operation.
- Users can stage the data themselves or can ask CAF administrator to do it for them.

Software packages distribution: alien (packman)

- ALICE analysis software (AliRoot) is released weekly in the form of so-called “analysis tags” (package name ends with “-AN”) that incorporate weekly developments and, less frequently, in the form of “revision tags” (package names end with “-Rev”) that incorporate developments for longer period, tested features.
- Once analysis package is released it is made available on grid. Within 15 hours the package is automatically installed on CAF, too.
- The installation is done with the use of *alien package manager, packman*.

What is CAF for?

- CAF is intended for prompt interactive parallel analysis of physics data
 - Interactive – You can see the output before job finishes working (contrary to batch system where user gets the output after job has finished its execution)
 - Parallel – The data to be analysed is divided into subsets and processed in parallel
 - Data analysis on CAF is transparent – Users run the as if they were running on laptop/desktop with very few modifications
 - No special registration is needed – If you are a member of ALICE Virtual Organisation and can use grid, you can can use CAF, too!

Comparison to local processing and grid

- Instead of 1, 2 or 4 cores (normally) user gets $N=58$. This speeds up the processing by N times *minus the time required for merging*. Also, if the data is pre-staged, files are accessed almost as if they were on local disk.
- CAF is not intended as a replacement or alternative to grid. Only selected data is copied to CAF storage from grid. CAF can be used for testing new code quickly with not very large statistics. On the other hand, grid is normally used for processing of very large amount of data with tested and stable code.

Monitoring

- Using the standard ALICE monitoring tool, MonALISA, one can monitor the following information about CAF
 - Statuses of nodes of the cluster (up, down)
 - Statuses of proof and xrootd processes (up, down)
 - Network consumption
 - Memory consumption by user tasks
 - Usage of the disk storage (used, free disk space)
- Available from:
 - <http://pcalimonitor.cern.ch> -> CAF monitoring -> Machine views
 - <http://aaf.cern.ch> -> Favourite menu -> ALICE monitoring -> CAF

ALICE Analysis Facilities (AAF)

- CAF is not the only analysis facility (AF) available to ALICE users. Following AFs, though less powerful than CAF, are used quite extensively by ALICE users:
 - SKAF (Slovak Analysis Facility, Slovakia)
 - KIAF (KISTI Analysis Facility, South Korea)
 - SAF (Subatech Analysis Facility, France)
 - LAF (Lyon Analysis Facility, France)
 - JRAF (JINR Analysis Facility, Russia)

Statistics on CAF

- In average, weekly:
 - 15-20 users
 - ~100 TB of processed data
 - 600-800 M events

Summary

- CAF is a resourceful cluster for prompt interactive parallel analysis of pp and Pb-Pb data in ALICE. It uses a bunch of software for parallel analysis of subsets of physics data, management of those data as well as keeping up-to-date with releases of ALICE analysis software.
- CAF is by no means a replacement for grid. Only selected amount of data is copied to CAF for analysis. CAF is suitable for quick checks of data samples and analysis code while grid jobs normally process huge amounts of data with stable code to collect large statistics.
- Using of CAF does not require special registration or heavy modifications of the code.

Documentation and user support

- AAF web site: <http://aaf.cern.ch>
 - Tutorial for AAF users
 - Documentation for cluster administrators
 - Links to monitoring information
- AAF savannah, where users and cluster administrators report bugs concerning AAF software, request new features. Here users can also request pre-staging of data to analysis.