PROOF plans for 5.28/00

G. Ganis

PROOF: current issues

- Layout and setup
  - XrdProofd instabilities (next slides)
  - Complete dynamic setup (add workers on-the-fly)
  - Worker auto-discovery
- PAR packages: new TPackageManager
  - Better handling/detection required by the dynamic setup
  - Portability: PARs not currently supported on Windows (move to zip)
- Performance studies
  - Packetizers optimization / re-writing
    - TPacketizerAdaptive does not seem to fit ALICE needs
- Benchmark suite (Sangsu Ryu)
- PROOF-Lite on Windows
- Multi-master mode
  - Dynamic setup
  - TDataSetManagerSM, TPacketizerSM
    - Handle distributed datasets, dynamic work distribution
  - Support for pure CPU tasks

XrdProofd instabilities

- Under heavy load, currently at CAF we have one restart per day due to (some) users unable to connect
- This is better than before, but backtraces (typically) show still residual deadlock situations
- Restarts have been automatized, but affect all users because XrdProofd is a single point of failure (see next slide), which is bad

- We think (Fons and I) that the plug-in can be re-designed so to remove the deadlock cases
- However, we believe that the first thing to do it is to reduce the dependency of the session from XrdProofd (next slide)

XrdProofd: proposed layout modification

- Currently XrdProofd is a single point of failure
- Clients start a session via XrdProofd and use it as proxy to the session

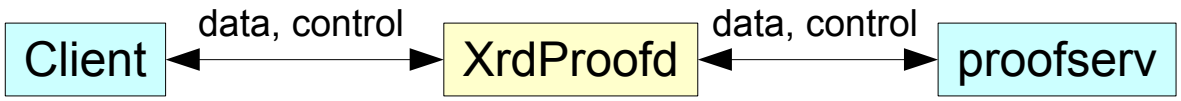

- The proposal is to keep the current connection for administration / control matters, but to have a direct TCP connection for the working (data) messages (TMessage)

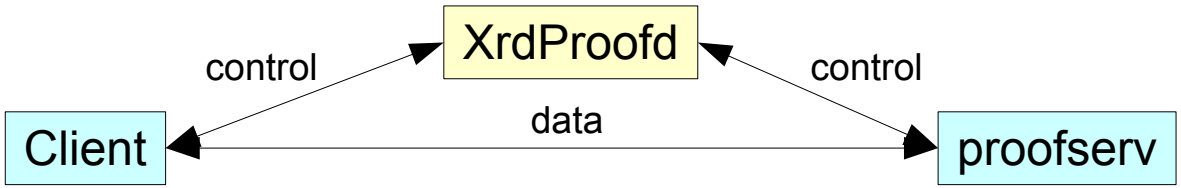

- This is achieved by adding a new TSocket member to TXSocket initiated by a callback of the session (proofserv) to the client
- Failure/freeze of XrdProofd would no more affect sessions
- Main issue: client firewall (flexible use of tunnels)

PROOF plans for 5.27/02

- First version of layout where sessions establish their own TCP connection by calling back the client/master
- Freeze re-design of new XrdProofd plug-in (w/ Fons)

- First version of the new benchmark suite (Sangsu Ryu)

- New features in dataset management
  - Possibility to run on multiple dataset at once
    - ALICE request; almost ready
  - Support for associated files
    - ALICE request; bulk ready (w/ Jan Fiete)

- PROOF-Lite for Windows (Bertrand; done already)
  - no PAR packages

PROOF plans for 5.27/04

- Consolidation of the new TCP layout
- First version of the re-designed XrdProofd plug-in

- Consolidation of the benchmark suite (Sangsu Ryu)

- First version of TPackageManager (Fons)

- First version of worker dynamic addition

PROOF plans / goals for 5.28/00

- Consolidation of the re-designed XrdProofd
- Consolidation of TPackageManager
- Completion and consolidation of the dynamic startup

Remarks

- The plans assume the new production release sometime in summer and one person working full time
  - They should be realistic, although tight

- Not listed in the plans:
  - Worker auto-discovery
  - TPacketizer improvements
  - Multi-master developments

- Some of these may (partially) come
- The multi-master developments are a key ingredient of the ALICE plans for PROOF on the Grid.

Remarks (2)

- Merging
  - Recently we have introduced
    - Automatic dataset creation for files created on the workers
      - Avoids hard merging; used by ALICE in the reconstruction
    - Parallel merging via submergers
      - Very useful, eg, when merging large numbers of histos
  - Hard file merging, i.e. the reduction of a large number of files to one or a few files, is not strictly a PROOF development
  - PROOF would definitely take advantage of a ROOT-level solution, which should be easy to integrate in the system

- Code coverage
  - Currently almost 0 because no test is run
  - Should enable stressProof in roottest/root/stress
  - New tests should be / are added regularly to stressProof