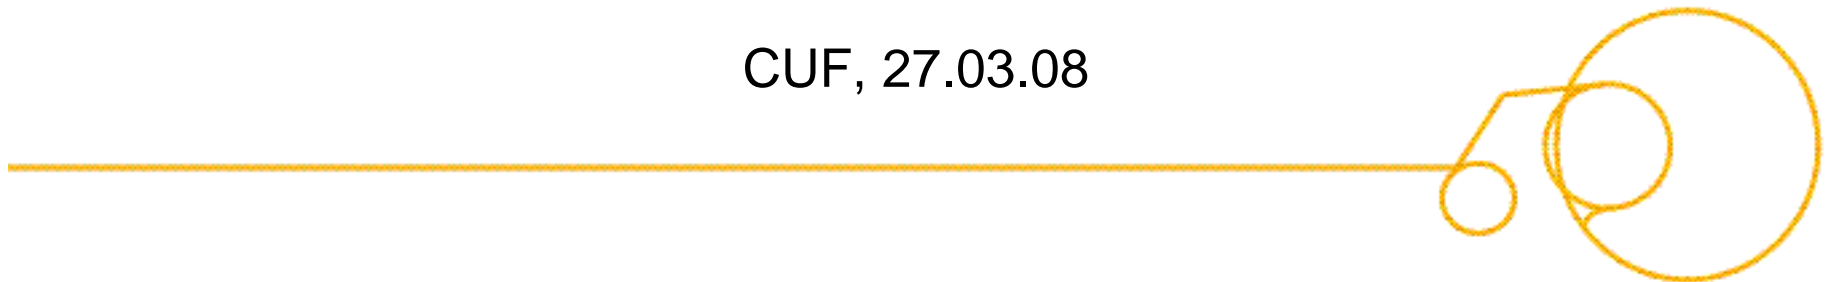


Good user practices + Dynamic staging to a CAF cluster

Jan Fiete Grosse-Oetringhaus, CERN PH/ALICE

CUF, 27.03.08






Good User Practices




- Before you start using CAF
 - Subscribe to alice-project-analysis-task-force@cern.ch using CERN SIMBA (<http://listboxservices.web.cern.ch/listboxservices>)
 - Read <http://aliceinfo.cern.ch/Offline/Analysis/CAF>
- Code development
 - Try your code on at least 2 files locally
 - 1 file may hide problems when switching to the next file
 - Run your code "as in PROOF"
 - Just change "proof" to "local" in StartAnalysis
 - Run "full PROOF"
- Don't use TProof::Reset if it is not needed (current issue)

CAF Schema



CAF computing cluster



Staging – Technical side



- Step 3 (now): Automatic
 - Staging script plugged into olbd
 - Implementation of PROOF datasets (by ALICE)
 - Staging daemon that runs on the cluster
 - Transparent migration from AliEn collection to PROOF datasets
 - Convenient for users, quota-enabled, garbage collection

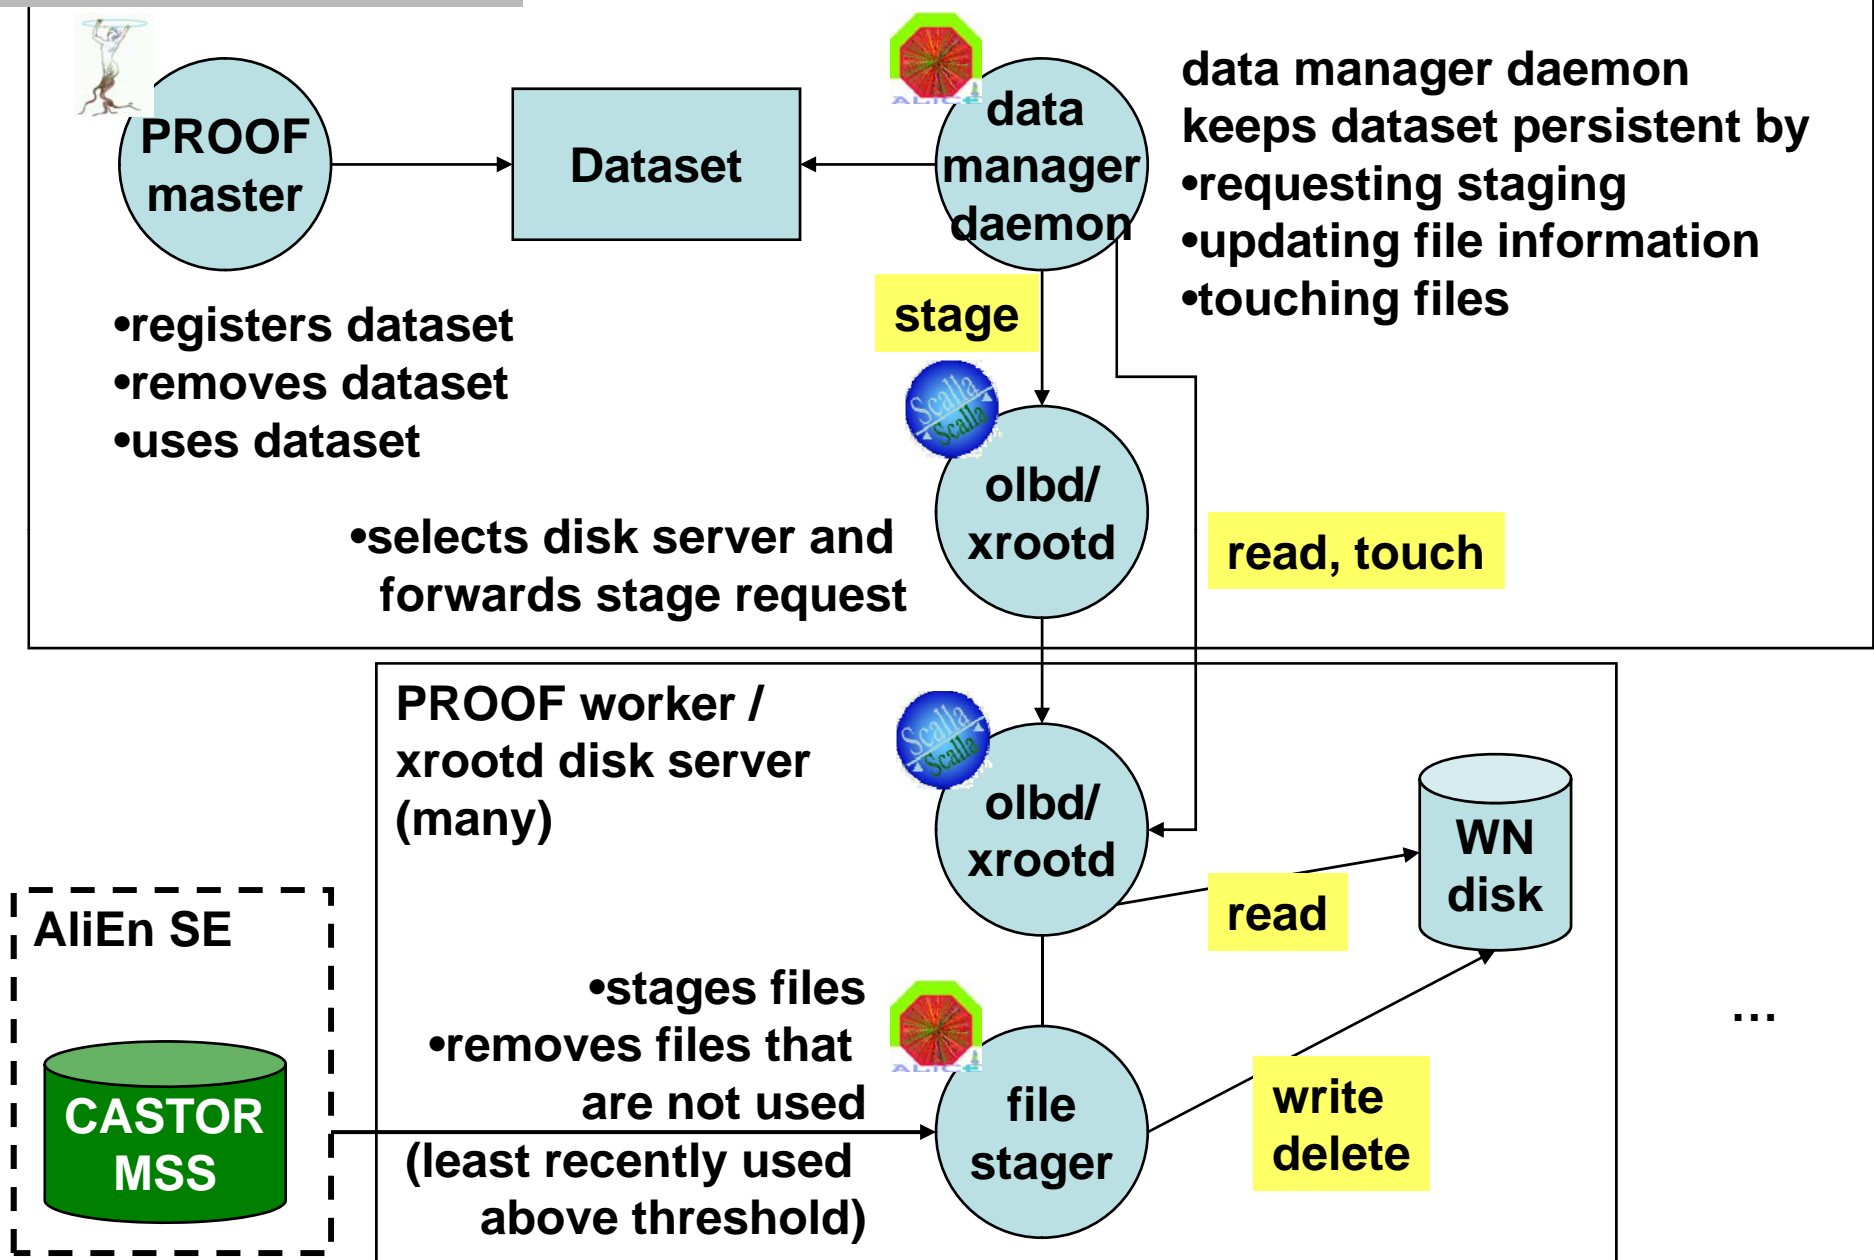
Introduction of PROOF datasets



- A dataset represents a list of files (e.g. physics run X)
 - Correspondence between AliEn collection and PROOF dataset
- Users register datasets
 - The files contained in a dataset are automatically staged from AliEn (and kept available)
 - Datasets are used for processing with PROOF
 - Contain all relevant information to start processing (location of files, abstract description of content of files)
- File-level storing by underlying xrootd infrastructure
- Datasets are public for reading (you can use datasets from anybody!)
- There are common datasets (for data of common interest)

Dataset concept

PROOF master / xrootd redirector



Staging script




- Two directories configured in xrootd/olbd for staging
 - /alien
 - /castor
- Staging script given with olb.prep directive
 - Perl script that consists of 3 threads
 - Front-End: Registers stage request
 - Back-End
 - Checks access privileges
 - Triggers migration from tape (CASTOR, AliEn)
 - Copies files, notifies xrootd
 - Garbage collector: Cleans up following policy file with low/high watermarks (least recently used above threshold)

Data manager daemon



- Keeps content of datasets persistent on disk
- Regularly loops over all datasets
- Sends staging requests for new files
- Extracts meta data from recently staged files
- Verifies that all files are still available on the cluster (by touch, prevents garbage collection)
 - Speed: 100 files / s

Datasets in Practice




- Create a AliEn collection (in aliensh)
 - `find -c myCollection /alice/sim/2007/LHC07c/pp_minbias/8051 root_archive.zip`
 - `find -c myCollection /alice/sim/2007/LHC07c/pp_minbias/8051 AliESDs.root`
- Use a ROOT version that supports datasets
 - LXPLUS: `source /afs/cern.ch/alice/caf/caf-lxplus-datasets.sh`
 - OR: Check out from ROOT SVN: `branches/dev/proof`
- Create DS from AliEn collection
 - Connect to AliEn
 - `TGrid::Connect("alien://")`
 - `gridColl = gGrid->OpenCollection("alien:///alice/cern.ch/user/j/jgrosseo/myCollection")`
 - `proofColl = gridColl->GetFileCollection();`
 - `proofColl->SetAnchor("AliESDs.root"); // collection of root_archive.zip`

Datasets in Practice (2)



- Upload to PROOF cluster
 - Connect to PROOF
 - TProof::Mgr("lxb6046")->SetROOTVersion("vPROOFDSMGR");
 - TProof::Open("lxb6046");
 - gProof->RegisterDataSet("myDataSet", proofColl);
- Check status
 - gProof->ShowDataSets();
- Use it
 - mgr->StartAnalysis("proof", "myDataSet");

Dataset in Practice (3)



- List available datasets: `gProof->ShowDataSets()`

```
root [3] gProof->ShowDataSets()
Dataset URI                                     |# Files|Default tree|# Events| Disk | Staged
/default/jgrosseo/ESD100                       | 6764|/esdTree | 676400| 343 GB| 100 %
/default/jgrosseo/run82XX_part1                | 10000|/esdTree | 998900| 288 GB| 99 %
/default/jgrosseo/run82XX_part2                | 10000|/esdTree | 944700| 272 GB| 94 %
/default/jgrosseo/run82XX_part3                | 10000|/esdTree | 987900| 285 GB| 98 %
/default/jgrosseo/ESD600                       | 1844|/esdTree | 184400| 51 GB| 100 %
/default/jgrosseo/ESD_FullMisalignment         | 944|/esdTree | 92100| 47 GB| 97 %
/default/jgrosseo/run12000                     | 62|/esdTree | 49| 4 GB| 79 %
/default/jgrosseo/ESD5000                      | 21899|/esdTree | 2189800| 1096 GB| 99 %
/default/jgrosseo/ProofSessionFiles            | 25438|/esdTree | 4460900| 1640 GB| 100 %
```

- You always see common datasets and datasets of your group