A Large Ion Collider Experiment

Analysis of data and storage quality of a distributed storage system

Adrian Eduard Negru (Politehnica University of Bucharest)
for the ALICE Collaboration

Motivation

- 120PB and 60 independent disk storage elements
- 2 billion individual files
- Data of different age, some 10 years old
- SE reliability must be kept at >95% level
- All files cannot be read continuously and SE elements (servers, disks) cannot be checked 100%
- Lost or degraded data must be found with sufficient accuracy and restored as soon as possible
- Storage element with time-increasing amount of degraded data or performance degradation must be flagged early for the sysadmins to take corrective action

File Catalogue overview and function

File catalogue content

- LFNs (unix-like) with ACLs
- PFNs (physical files, linked to SE)
- MD5sum, metadata
- Size, date, etc...

- Is the authoritative source of information for the entire processing chain, from raw data to end-user analysis
- Contains ***all*** files stored on ***all*** SEs from the beginning of times
- Allows for systematic analysis of data and storage quality without other sources of information

File crawler overview

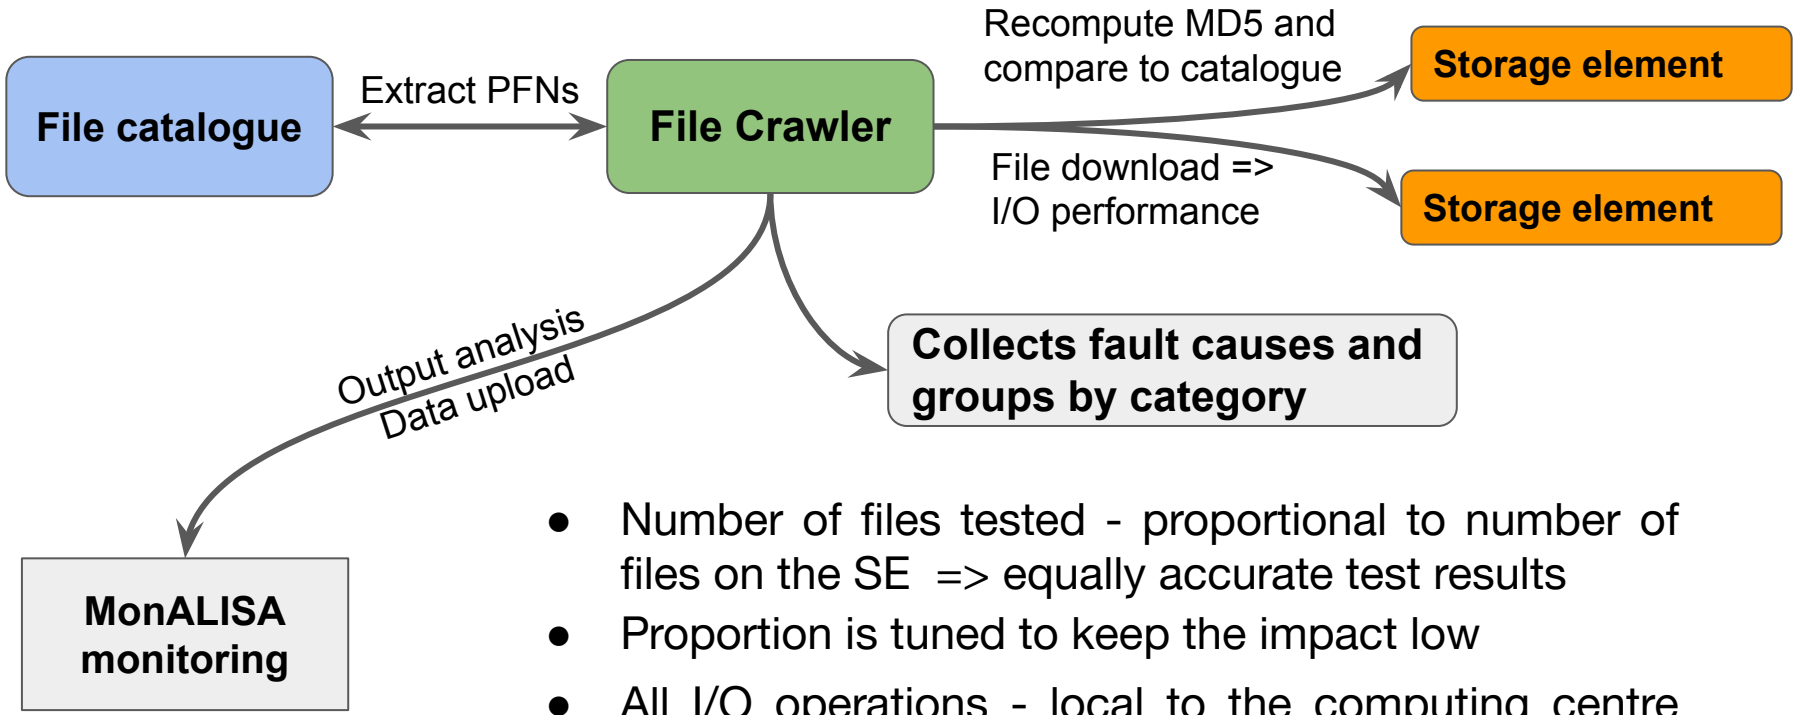

- Number of files tested - proportional to number of files on the SE => equally accurate test results
- Proportion is tuned to keep the impact low
- All I/O operations - local to the computing centre (same as Grid payloads)

Architecture

- Periodically launch crawling jobs to the Grid
- Jobs are split according to data location (per SE)
- Completed crawling for an SE triggers analysis and subsequent corrective action + results are inserted into monitoring database
- The workflow below is applied in parallel for all SEs

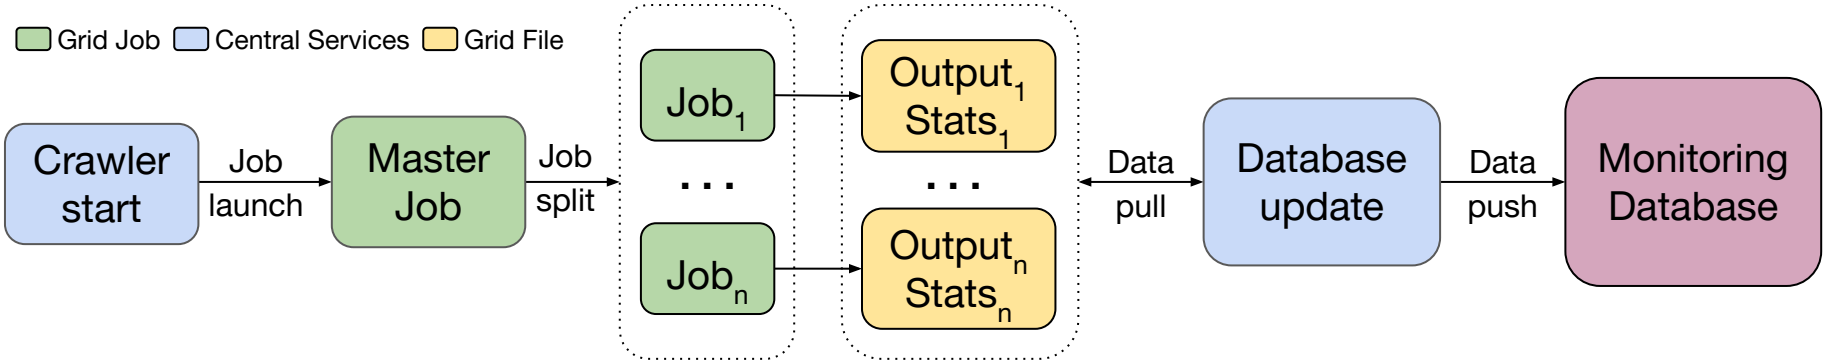

Sample size calculation

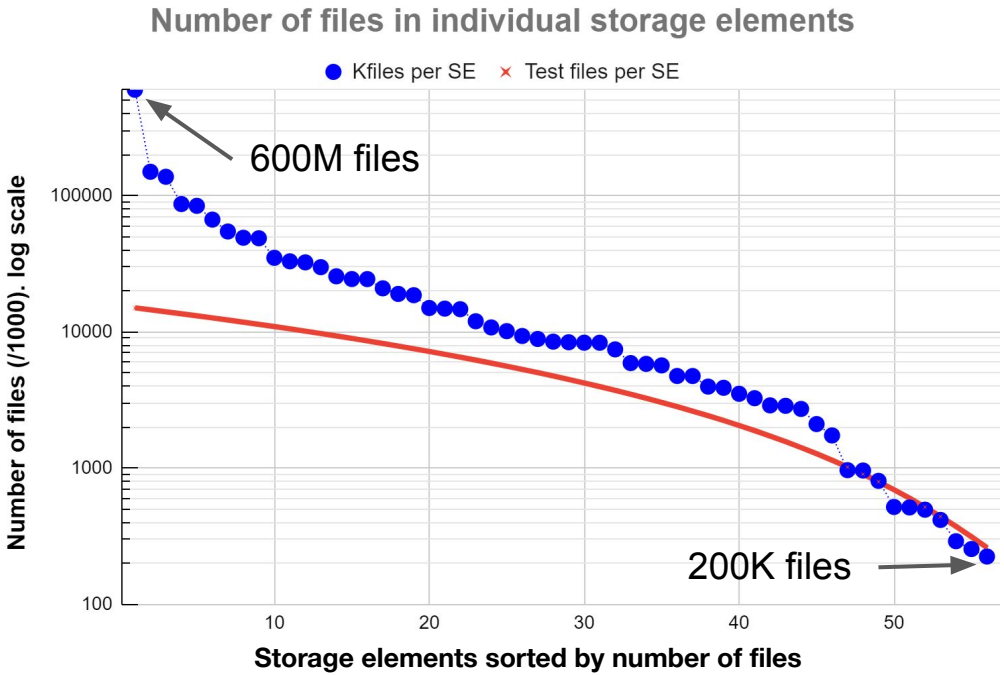

Test sample tuned to minimize the load on the SEs while ensuring accurate and meaningful test sample

The number of tests vary from ~0.001% of the total at the high end to ~100% at the low end

Analysis of status codes

- Up to 3% of the total sample size fault detected in certain periods!
- SE response depends on load - crawler to be tuned to detect this

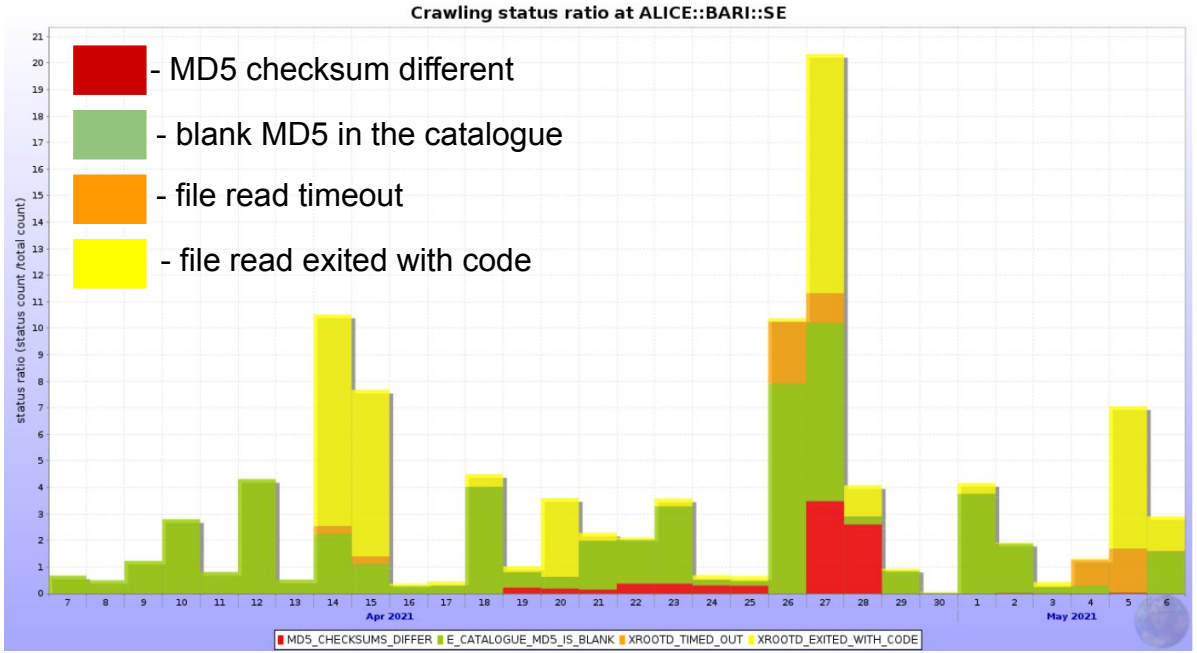

- Missing MD5 in catalogue - retrofitting by first analysing the file through other methods
- Low percentage of really corrupted data

Throughput analysis

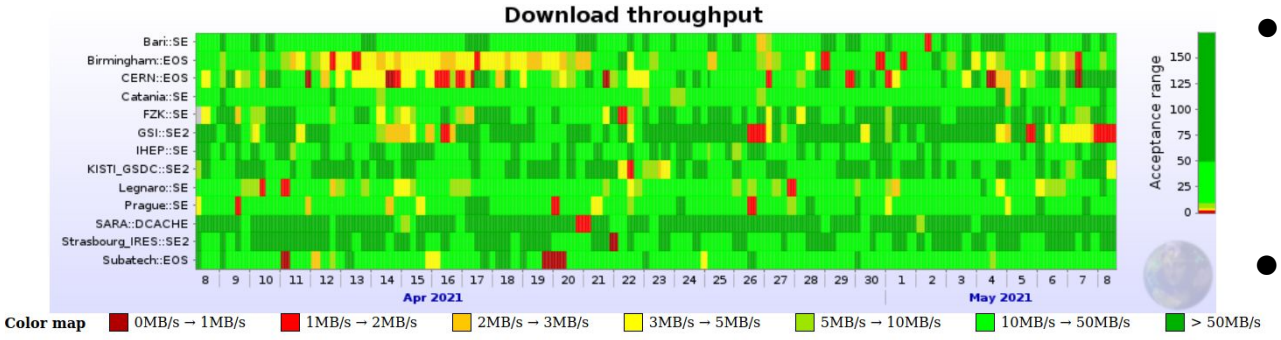

Statistics

Link name	Data		Average
	Starts	Ends	Download throughput
Bari::SE	08 Apr 2021 05:14	08 May 2021 09:23	39.97 MB/s
Birmingham::EOS	08 Apr 2021 05:15	08 May 2021 09:21	14.68 MB/s
CERN::EOS	08 Apr 2021 05:19	08 May 2021 09:43	27.37 MB/s
Catania::SE	07 Apr 2021 23:03	08 May 2021 09:25	26.80 MB/s
FZK::SE	08 Apr 2021 11:25	08 May 2021 09:21	45.23 MB/s
GSI::SE2	08 Apr 2021 05:13	08 May 2021 09:14	58.79 MB/s
IHEP::SE	08 Apr 2021 05:15	08 May 2021 09:23	44.28 MB/s
KISTI_GSDC::SE2	08 Apr 2021 05:17	08 May 2021 09:08	46.00 MB/s
Legnaro::SE	08 Apr 2021 05:16	08 May 2021 09:06	23.10 MB/s
Prague::SE	08 Apr 2021 05:17	08 May 2021 09:18	26.75 MB/s
SARA::DCACHE	08 Apr 2021 05:18	08 May 2021 09:19	95.15 MB/s
Strasbourg_IRES::SE2	08 Apr 2021 05:16	08 May 2021 09:10	49.27 MB/s
Subatech::EOS	08 Apr 2021 05:15	08 May 2021 09:18	32.43 MB/s

- Average throughput per client for multiple crawler iterations (plot shows 1 month of data)
- and - below optimal (5MB/s) throughput => typical reason is heavy analysis load
- Collected data to be used to optimize analysis patterns

Analysis of the samples

- Examine samples inspected by the crawler
  - Analyse the error message
  - Reproduce the error by re-executing the command that failed
- In the future - debugging tips and possible fixes to common issues

		MD5_CHECKSUMS_DIFFER,XRDFS_C			Filter
PFN	**SE Name**	**Status Name**	**Size (B)**	**Timestamp**	**Status Message**
root://neos.nipne.ro:1094//08/26705/b54354f8-1227-11e8-8c69-2ba5bf90fcb5	ALICE::NIHAM::EOS	MD5_CHECKSUMS_DIFFER	98401595	Apr 22,2021 06:49	Local file doesn\'t match catalogue details MD5 c...
root://grinr03.inr.troitsk.ru:1094//13/61373/98f6d712-42ee-11e0-8740-0019bbc6247c	ALICE::TROITSK::SE	MD5_CHECKSUMS_DIFFER	197750808	Apr 23,2021 09:35	Local file doesn\'t match catalogue details MD5 c...
root://ali-nxrd.to.infn.it:1094//03/59050/73e90802-757d-11e0-8a17-b7d2e3473590	ALICE::TORINO::SE2	MD5_CHECKSUMS_DIFFER	53057140	Apr 14,2021 19:08	Local file doesn\'t match catalogue details MD5 c...
root://alice-t1-eos-mgm01.sdfarm.kr:1094//03/00517/358bab02-848d-11e7-b871-cfae48f7be88	ALICE::KISTI_GSDC::EOS	MD5_CHECKSUMS_DIFFER	130366274	Apr 19,2021 17:59	Local file doesn\'t match catalogue details MD5 c...
root://neos.nipne.ro:1094//08/21339/846c57cc-c035-11e9-b56a-27fe2bbbf9cc	ALICE::NIHAM::EOS	MD5_CHECKSUMS_DIFFER	12211	Apr 14,2021 19:04	Local file doesn\'t match catalogue details MD5 c...
root://clralicexrd.in2p3.fr:1094//04/12325/a19824f6-f1c6-11e6-ba09-ebd157cb4b4a	ALICE::CLERMONT::SE	XRDFS_CANNOT_CONFIRM_UPLOAD	279128250	Apr 14,2021 19:06	[/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Package...
root://clralicexrd.in2p3.fr:1094//01/29766/831e1676-f1c5-11e6-b0df-138910a3cca0	ALICE::CLERMONT::SE	XRDFS_CANNOT_CONFIRM_UPLOAD	87568203	Apr 14,2021 19:07	[/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Package...
root://clralicexrd.in2p3.fr:1094//04/55516/b9b4994c-efd3-11e6-8e73-d76fe48573a0	ALICE::CLERMONT::SE	XRDFS_CANNOT_CONFIRM_UPLOAD	3274	Apr 14,2021 19:08	[/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Package...
root://clralicexrd.in2p3.fr:1094//04/35439/2d603a12-f1c4-11e6-8b5a-b3cae7df7731	ALICE::CLERMONT::SE	XRDFS_CANNOT_CONFIRM_UPLOAD	35667565	Apr 14,2021 19:09	[/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Package...
root://clralicexrd.in2p3.fr:1094//01/16208/89beaa56-ed4f-11e6-b979-9b28e76b708d	ALICE::CLERMONT::SE	XRDFS_CANNOT_CONFIRM_UPLOAD	554273	Apr 14,2021 19:07	[/cvmfs/alice.cern.ch/x86_64-2.6-gnu-4.1.2/Package...

ALICE

Summary

- A file crawler system determines the health and performance of the ALICE Storage Elements
  - Analyses a representative data subset in each iteration
  - Detects corrupted files
  - Splits errors into categories for easy debug
  - Analyses the storage element load
- Data is visualized in MonALISA
- As of May 2021, more than 11 million files analysed
- Future work
  - More fine-grained performance analysis
  - Early fault detection
  - Suggest action for sysadmins and heavy users