JAliEn status

ALICE-USA meeting
LBL, 2025-09-17

costin.grigoras@cern.ch

Deployment status

	1 year ago	now
JAliEn	1.9.2	2.0.3
Whole node queues	12	15 (+2 EPN)
16-core queues	0	7
1 / 2 / 3 day TTL	all / 0 / 0	63 / 2 / 6
EL 7 / 8 / 9	9% / 29% / 62%	5% / 5% / 90%
AVX2	96% nodes (426 nodes w/o)	95% nodes, 98.6% cores
ORNL		54 nodes / 40 nodes
LBL / Perlmutter		All ok (?)
AVX512	24% nodes	24.5% nodes (and cores)

JAliEn highlights

Exploit cgroups v2 availability

- Not only resource constraining
- Replacing `/proc/` or `` `ps` ``-based monitoring
  - CPU and memory usage, number of processes and threads
- cgroups v2 is now a requirement

Recording all job resubmissions

- Currently only the last attempt is visible in `` `ps` ``
- Todo: new command to interface with the archive table

TTL optimization

Enabled by default for all productions

- Learning for the first 24h of a production
- Can be disabled explicitly (`TTLOptimizationType="false"`);

Further reduced for IO jobs

- Scaling with the number of input files
- Or with the volume of data they can process

Good results, contributing to better resource usage

- Looking into doing the same to learn actual disk space requirement from jobs (`WorkDirectorySize`)

Apptainer -> Podman

Recent development, prompted by large EXPIRED job rate at Kisti

Traced by the site admins to system calls issued by the payloads

```
kill -9 0
```

Apptainer does not isolate the process group thus the signal was delivered to all jobs in the slot

Podman solves this issue and is available by default in the OS

- Requires separate images, in `/cvmfs/alice.cern.ch/containers/fs/podman/`
- Complicates monitoring, as payload is no longer a direct descendent of the `JobWrapper` process

GPU brokering

Advertise WNs capabilities ✅

- Count of *(GPU vendor, model, architecture)* tuples

Reserve and pass them down to the containers ✅

JDL tags to come

- Jobs are to advertise all the GPU arch compiled in the binaries
- Matching should do the intersection with the WN hardware
- Also matching the number of cards (granularity is an entire card)
- Can build on the regex constraints

GPU usage monitoring

…

Job kill reasons

Added numeric subcodes to ERROR_E jobs

Available in command line `ps`, with decoding (-v)

```
jsh:[alice] [29] /alice/cern.ch/user/g/grigoras/   > ps -f f -a -X -s LBL,ORNL -t 1,0 -H -vv -S -N
       Owner           PID   Masterjob          Queue name    Worker node  STS   ErrNo     Reason  CPUs                                Command
     aliprod    3412072249  3412067554    ALICE::ORNL::ORNL    alice-125   EE     143  IDLE_CPU      8          o2_o2dpg_sim_anchored_dev2.sh
     aliprod    3412072267  3412067554    ALICE::ORNL::ORNL    alice-114   EE     143         OK     8          o2_o2dpg_sim_anchored_dev2.sh
    alitrain    3412509804  3412509728    ALICE::ORNL::ORNL    alice-006   EE       6         OK     1                        lego_train.sh
alihyperloop    3412519694  3412519678    ALICE::ORNL::ORNL    alice-112   EE     137         OK     1                       run_train12.sh
```

And also as monitoring stream

Idle CPU, Over TTL or Over memory

CERN/ALIEN_KILL_REASON_Nodes_Summary/sum/<user>_<code>

Job summaries in `ps`

```
Summary of sites
ALICE::ORNL::ORNL                       :     290 jobs (    688 cores) (96.03%)
ALICE::LBL::Perlmutter                  :      12 jobs (     30 cores) (3.974%)

Summary of worker nodes
alice-126                               :      16 jobs (     42 cores) (5.298%)
alice-044                               :      14 jobs (     27 cores) (4.636%)

Summary of users
      aliprod :     205 jobs (    602 cores) (67.88%)
 alihyperloop :      38 jobs (     38 cores) (12.58%)
     alitrain :      28 jobs (     28 cores) (9.272%)

Summary of process exit codes (for error jobs)
          1 :     184 jobs (    371 cores) (60.93%) (Snapshot (?))
        137 :      86 jobs (    208 cores) (28.48%) (ROOT caught a KILL signal (?))
        143 :      14 jobs (    112 cores) (4.636%) (ROOT caught a TERM signal (?))

Summary of kill reasons (when JAliEn killed them explicitly)
          0 (      Finished OK or no kill decision taken) :     290 jobs (    622 cores) (96.03%)
         50 (    Payload idling for more than 15 minutes) :      12 jobs (     96 cores) (3.974%)

Summary of states
     ERROR_E ( -3) :     288 jobs (    681 cores) (95.36%)
     ERROR_V (-10) :      11 jobs (     25 cores) (3.642%)
     EXPIRED (-12) :       3 jobs (     12 cores) (0.993%)
```

`ps` in alimonitor

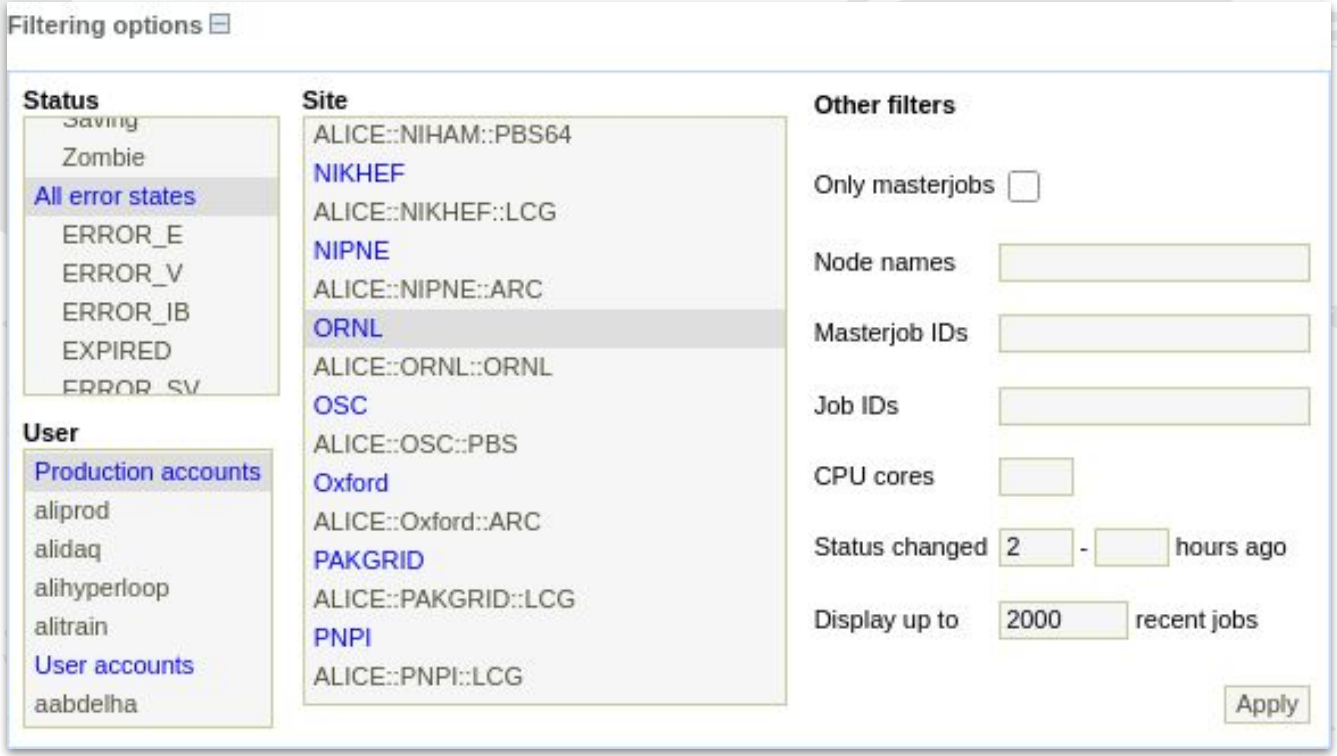

aliprod	3412529245	3412529056	ALICE::ORNL::ORNL	alice-029	ERROR_E	Snapshot (1)	OK (0)	2	535.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 4s	20.31%
aliprod	3412529247	3412529056	ALICE::ORNL::ORNL	alice-022	ERROR_E	Snapshot (1)	OK (0)	2	871 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:37	2m 5s	33.06%
aliprod	3412529249	3412529056	ALICE::ORNL::ORNL	alice-113	ERROR_E	Snapshot (1)	OK (0)	2	198.8 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 18s	7.547%
aliprod	3412529250	3412529056	ALICE::ORNL::ORNL	alice-132	ERROR_E	Snapshot (1)	OK (0)	2	169.4 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:37	1m 35s	8.889%
aliprod	3412529251	3412529056	ALICE::ORNL::ORNL	alice-122	ERROR_E	Snapshot (1)	OK (0)	2	455.2 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 4s	9.375%
aliprod	3412529253	3412529056	ALICE::ORNL::ORNL	alice-cray017	ERROR_E	Snapshot (1)	OK (0)	2	221.5 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	39s	16.67%
aliprod	3412529254	3412529056	ALICE::ORNL::ORNL	alice-122	ERROR_E	Snapshot (1)	OK (0)	2	502.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	34s	17.65%
aliprod	3412529255	3412529056	ALICE::ORNL::ORNL	alice-044	ERROR_E	Snapshot (1)	OK (0)	2	3.395 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 4s	64.84%
aliprod	3412529257	3412529056	ALICE::ORNL::ORNL	alice-022	ERROR_E	Snapshot (1)	OK (0)	2	1.17 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 5s	35.38%
aliprod	3412529258	3412529056	ALICE::ORNL::ORNL	alice-128	ERROR_E	Snapshot (1)	OK (0)	2	177.1 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 4s	9.375%
aliprod	3412529260	3412529056	ALICE::ORNL::ORNL	alice-022	ERROR_E	Snapshot (1)	OK (0)	2	470.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 5s	19.23%
aliprod	3412529261	3412529056	ALICE::ORNL::ORNL	alice-126	ERROR_E	Snapshot (1)	OK (0)	2	166.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	34s	8.824%
aliprod	3412529262	3412529056	ALICE::ORNL::ORNL	alice-132	ERROR_E	Snapshot (1)	OK (0)	2	184.6 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:37	2m 8s	5.078%
aliprod	3412529263	3412529056	ALICE::ORNL::ORNL	alice-cray014	ERROR_E	Snapshot (1)	OK (0)	2	627.6 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:37	2m 1s	8.642%
aliprod	3412529264	3412529056	ALICE::ORNL::ORNL	alice-128	ERROR_E	Snapshot (1)	OK (0)	2	3.653 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 5s	42.31%
aliprod	3412529265	3412529056	ALICE::ORNL::ORNL	alice-056	ERROR_E	Snapshot (1)	OK (0)	2	3.672 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	57.03%
aliprod	3412529267	3412529056	ALICE::ORNL::ORNL	alice-023	ERROR_E	Snapshot (1)	OK (0)	2	215.5 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	1m 4s	17.05%
aliprod	3412529271	3412529056	ALICE::ORNL::ORNL	alice-132	ERROR_E	Snapshot (1)	OK (0)	2	170.4 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 7s	8.889%
aliprod	3412529273	3412529056	ALICE::ORNL::ORNL	alice-056	ERROR_E	Snapshot (1)	OK (0)	2	191.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	19.12%
aliprod	3412529276	3412529056	ALICE::ORNL::ORNL	alice-126	ERROR_E	Snapshot (1)	OK (0)	2	156.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:37	1m 4s	6.25%
aliprod	3412529278	3412529056	ALICE::ORNL::ORNL	alice-126	ERROR_E	Snapshot (1)	OK (0)	2	583 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:36	34s	20.29%
aliprod	3412529279	3412529056	ALICE::ORNL::ORNL	alice-119	ERROR_E	Snapshot (1)	OK (0)	2	351.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:47	1m 20s	10%
aliprod	3412529280	3412529056	ALICE::ORNL::ORNL	alice-cray014	ERROR_E	Snapshot (1)	OK (0)	2	517.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:38	2m 15s	7.407%
aliprod	3412529281	3412529056	ALICE::ORNL::ORNL	alice-cray017	ERROR_E	Snapshot (1)	OK (0)	2	3.759 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	48.84%
aliprod	3412529282	3412529056	ALICE::ORNL::ORNL	alice-044	ERROR_E	Snapshot (1)	OK (0)	2	384.4 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	35.29%
aliprod	3412529284	3412529056	ALICE::ORNL::ORNL	alice-128	ERROR_E	Snapshot (1)	OK (0)	2	163.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	8.594%
aliprod	3412529285	3412529056	ALICE::ORNL::ORNL	alice-021	ERROR_E	Snapshot (1)	OK (0)	2	458.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	18.75%
aliprod	3412529286	3412529056	ALICE::ORNL::ORNL	alice-019	ERROR_E	Snapshot (1)	OK (0)	2	899.8 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 5s	24.62%
aliprod	3412529287	3412529056	ALICE::ORNL::ORNL	alice-cray016	ERROR_E	Snapshot (1)	OK (0)	2	157.1 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	17.65%
aliprod	3412529288	3412529056	ALICE::ORNL::ORNL	alice-044	ERROR_E	Snapshot (1)	OK (0)	2	3.672 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:47	2m 6s	36.11%
aliprod	3412529290	3412529056	ALICE::ORNL::ORNL	alice-128	ERROR_E	Snapshot (1)	OK (0)	2	171.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	16.18%
aliprod	3412529291	3412529056	ALICE::ORNL::ORNL	alice-022	ERROR_E	Snapshot (1)	OK (0)	2	3.429 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	68.22%
aliprod	3412529292	3412529056	ALICE::ORNL::ORNL	alice-cray017	ERROR_E	Snapshot (1)	OK (0)	2	632.5 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 4s	17.05%
aliprod	3412529293	3412529056	ALICE::ORNL::ORNL	alice-044	ERROR_E	Snapshot (1)	OK (0)	2	172.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	1m 5s	11.45%
aliprod	3412529295	3412529056	ALICE::ORNL::ORNL	alice-056	ERROR_E	Snapshot (1)	OK (0)	2	154.3 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	17.65%
aliprod	3412529296	3412529056	ALICE::ORNL::ORNL	alice-cray011	ERROR_E	Snapshot (1)	OK (0)	2	207.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:47	2m 9s	6.202%
aliprod	3412529297	3412529056	ALICE::ORNL::ORNL	alice-cray014	ERROR_E	Snapshot (1)	OK (0)	2	802.7 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:38	2m 12s	10.98%
aliprod	3412529299	3412529056	ALICE::ORNL::ORNL	alice-039	ERROR_E	Snapshot (1)	OK (0)	2	3.328 GB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	111.8%
aliprod	3412529304	3412529056	ALICE::ORNL::ORNL	alice-054	ERROR_E	Snapshot (1)	OK (0)	2	449.8 MB	/alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh	today 11:46	34s	35.29%
3 entries	**306 jobs**		**1 entries**	**58 entries**	**2 entries**	**4 entries**	**2 entries**	**3 entries**	**2.657 GB**	**14 entries**		**41m 55s**	**36.11%**

Peak memory usage
Peak memory usage
Number of jobs
Peak memory allocation (MB)
Job memory usage
Average peak memory usage: 2720MB
Job duration histogram
Wall time per job
Number of jobs
Wall time (hours)
Job duration
Average job duration: 41m 55s
CPU efficiency histogram
CPU efficiency
Number of jobs
CPU efficiency (CPU time / wall time) (%)
Job efficiency
Average CPU efficiency: 36.11%
Breakdown of executed commands
Count Key
162 /alice/cern.ch/user/a/alidaq/AOD/AODAnalysisQC/run_AnalysisQC.sh
37 /alice/cern.ch/user/a/alihyperloop/bin/run_train12.sh
1 /alice/cern.ch/user/a/aliprod/EP24I1/runme.sh
1 /alice/cern.ch/user/a/aliprod/LHC22k4_nightly/o2_o2dpg_sim_v4.sh
54 /alice/cern.ch/user/a/aliprod/LHC23d1/o2_o2dpg_sim_anchored_dev2.sh
1 /alice/cern.ch/user/a/alitrain/PWGGA/GA_pp_AOD/3074_20250912-1457_child_13/lego_train.sh
Exit code (in case of error)
Count Key Description
173 1 Snapshot
5 6 extraValidation failed
89 137 ROOT caught a KILL signal
39 143 ROOT caught a TERM signal
Job status
Count Key
295 ERROR_E
11 ERROR_V
names
Count Key
1 alice-003
1 alice-004
1 alice-006
1 alice-007
1 alice-008
1 alice-009
1 alice-010
1 alice-012
4 alice-014
1 alice-015
4 alice-019
3 entries 306 jobs 1 entries 58 entries 2 entries 4 entries 2 entries 3 entries 2.657 GB 14 entries 41m 55s 36.11%

MonALISA client changes

New communication channel

- WebSocketS based, JSON serialization
- Simpler to interface with external tools
- Will reuse the VoBox host certificates

New graphical interface

- Web based (using the same *wss* endpoint for data)
  - Java WebStart on client machines no longer needed
- The proxy serves the interface and routes subscriptions/data
- Each service could also run one

Search...

- ALICE
  - alissandra01.cern.ch
    - IPs
    - Machine
    - Master
      - alissandra01.cern.ch
      - localhost
    - MonaLisa
    - MonaLisa_LocalSysMon
    - MonaLisa_ThPStat
    - UPS
  - alissandra02.cern.ch
    - IPs
    - Machine
      - CCISS1
      - System
      - disk
      - ipmi
      - kernel
      - memory
      - sensors
    - Master
    - MonaLisa
    - MonaLisa_LocalSysMon
    - MonaLisa_ThPStat

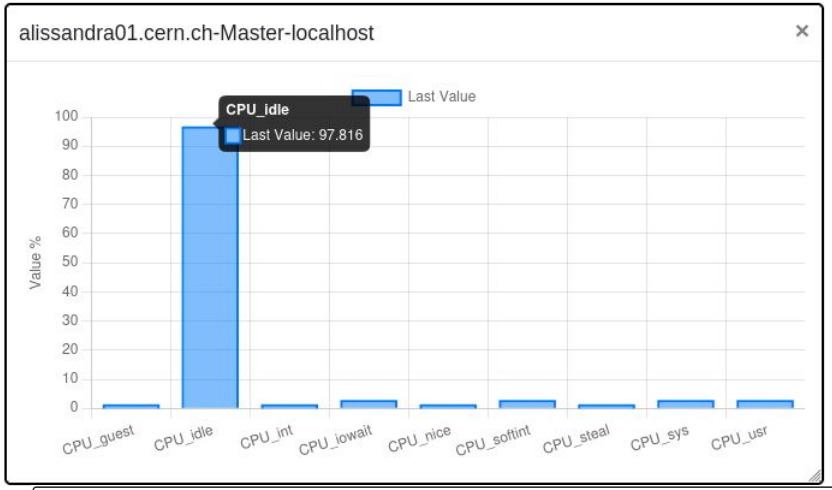

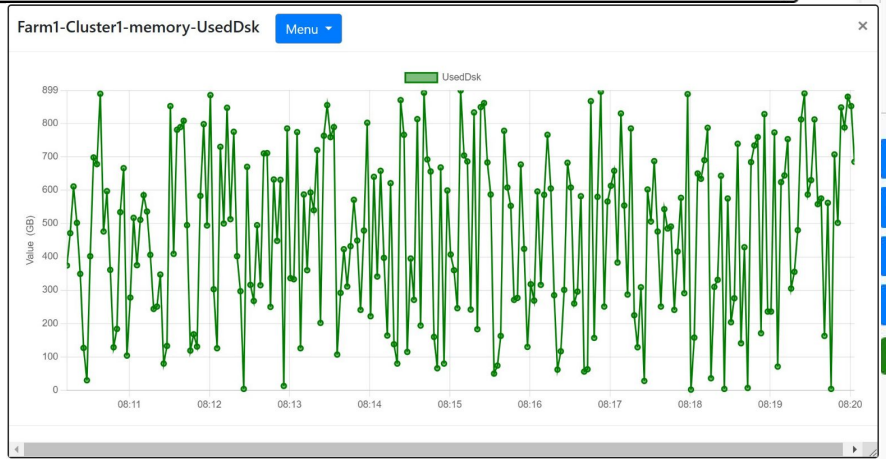

Parameters List

- br-5f3db5232a3d_IN
- br-5f3db5232a3d_OUT
- CPU_guest
- CPU_idle
- CPU_int
- CPU_iowait
- CPU_nice
- CPU_softint
- CPU_steal
- CPU_sys
- CPU_usr
- docker0_COLLS
- docker0_ERRS
- docker0_IN

History Chart

Last Value Chart

Global Zoom

Global Zoom Time Ago

Select all parameters

ALICE

IPv6 support

SiteSonar WN probe

- 77% of the WNs are ok
- 12% cannot connect - not dual stacked
- 6% time out / transfer errors
- 5% can't resolve IPv6

All components are IPv6 ready

- Java, Python, Xrootd 5+ (client and server)

IPv4 still required

- And as fallback in case IPv6 doesn't work
- Online farm is (and will stay) IPv4-only
- Legacy binaries compatibility

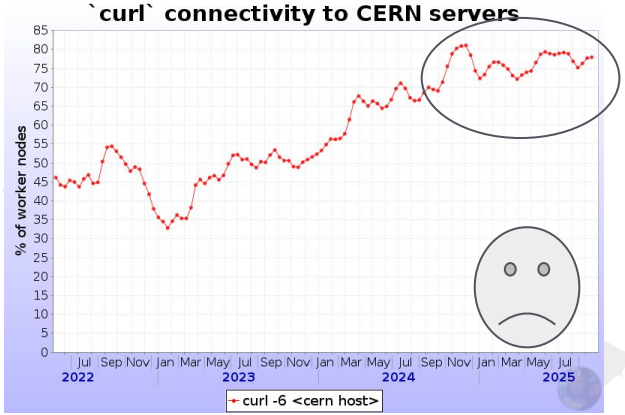

60% of the VoBoxes are dual stacked

89% of the storage volume is dual stacked & working

- 7 sites still don't have it (ORNL included), others show various IPv6 errors

Plans for “IPv6-only” WLCG data transfers For discussion with all stakeholders

LHCOPN (the easiest - just Tier 0/1)

- By end Run 3 encourage the deployment of IPv6 on *all* WLCG services (today ~80%)
- Need to continue to identify/remove IPv4 on LHCOPN
- Propose IPv4 peering removed from LHCOPN as soon as possible/sensible
- Before Data Challenge DC27 (should test what we plan to use) - my personal view
- See next talk (Martelli/McKee) - USA test removal of IPv4 on LHCOPN

For transfers over LHCONE:

- Identify/remove ongoing use of IPv4 on LHCONE network (WLCG traffic)
- Move the WLCG LHCONE to be IPv6-only (if confirmed to be sensible)

End point of the IPv6 work:

- Well before start HL-LHC Run 4 - exact date still to be agreed with the WLCG MB

21

FNAL

- FNAL stopped all LHCOPN IPv4 peerings at 15:53 UTC

BNL

- BNL stopped all LHCOPN IPv4 peerings at 11:27 UTC (7:27 AM EDT or 13:27 CERN time where plots were copied)

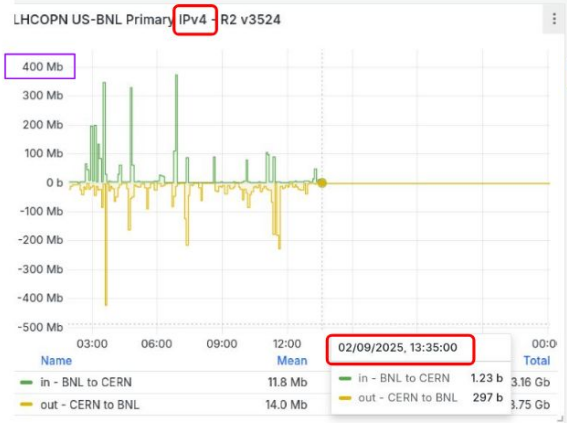

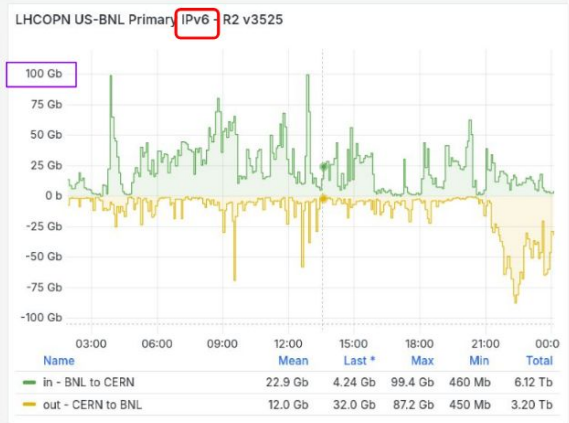

https://monit-grafana-open.cern.ch/d/Clqtof5nk/us-bnl?orgId=16

Data and ops over http

EOS features

- `fsck` for reporting and repairing
- HTTP endpoint for data access
- Access to `fsck` reports over the same http port
  - Unprivileged account with access just to this list
- Available from 5.2+
- See Andreea's talk at the EOS workshop

fsck report and repair

Set *scaninterval* for space and all filesystems

```
space config <space-name>(ex:default) space.scaninterval=<sec>
fs config <fs-id> scaninterval=<sec>
```

Activate collection or repair threads

```
fsck config toggle-collect [<threads_number>]
fsck config toggle-repair [<threads_number>]
```

HTTP(s) data and REST API

```
xrd.protocol XrdHttp:1094 /usr/lib64/libXrdHttp.so
http.exthandler EosMgmHttp /usr/lib64/libEosMgmHttp.so
eos::mgm::http::redirect-to-https=1
xrd.tls  /etc/grid-security/daemon/hostcert.pem /etc/grid-security/daemon/hostkey.pem
xrd.tlsca  certdir /etc/grid-security/certificates/
http.gridmap /etc/grid-security/grid-mapfile
EOS_MGM_ENABLE_REST_API=1
```

(in /etc/sysconfig/eos_env)

Xrootd and HTTP can run on the same port, no need to set up a different firewall for it

Storage balancing

Important when deploying new capacity

```
eos space config default space.balancer=on
eos geosched disabled add <node_geo_tag> plct \*
```

(wait to reach ~= level)

```
eos geosched disabled rm <node_geo_tag> plct \*
```

If already in hold/cold mix of storage nodes

Take turns draining each node to ~half

```
eos fs ls <node_name> --io
```

Might benefit from group-drainer

Xrootd client library

ALICE

- Implement the Xrootd protocol in Java
- Starting from the dCache code
  - Limited client support, mainly protocol definition
- No platform binaries as dependencies to Java
- Persistent connections, better error handling
  - Retry IO operations on SEs with replication factor >1
- Started with connection-intensive operations
  - `ls -lR`, `stat`, next `prepare`, `rm` and then the read/write ops

Job mix

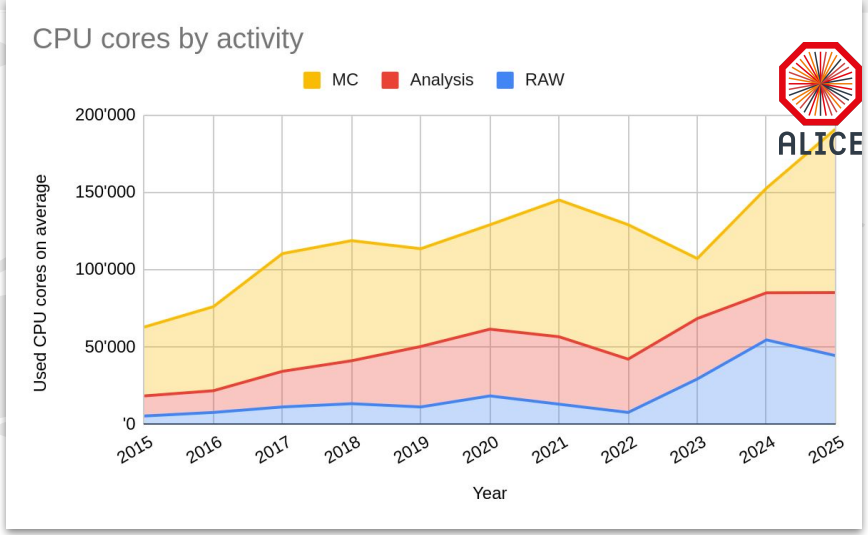
CPU cores by activity
MC
Analysis
RAW
Used CPU cores on average
200'000
150'000
100'000
50'000
'0
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
Year

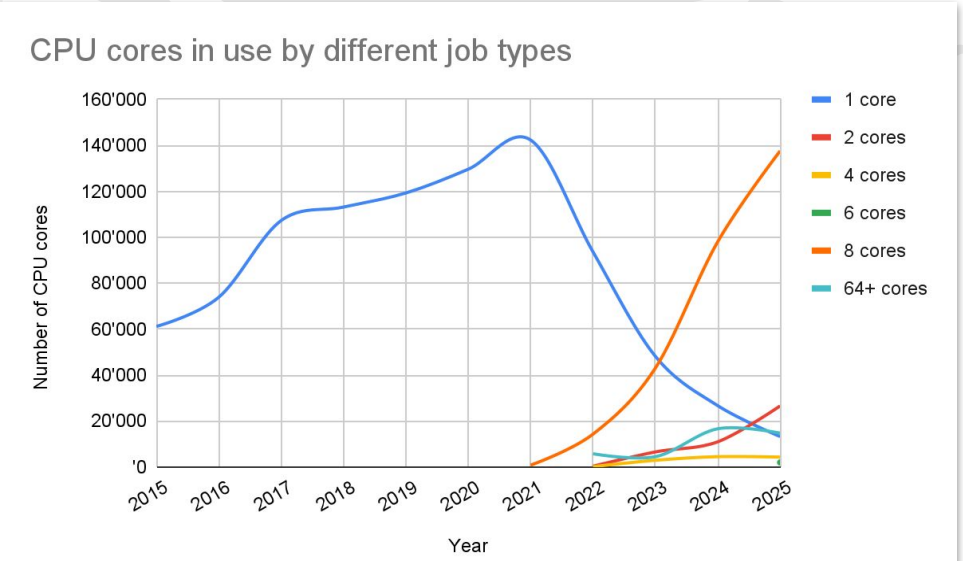
CPU cores in use by different job types
Number of CPU cores
160'000
140'000
120'000
100'000
80'000
60'000
40'000
20'000
'0
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
Year
1 core
2 cores
4 cores
6 cores
8 cores
64+ cores

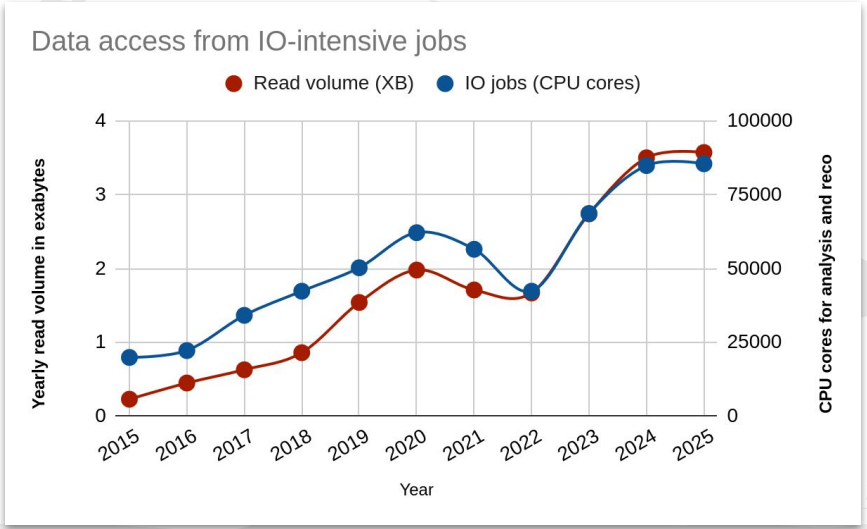
Data access from IO-intensive jobs
Read volume (XB)
IO jobs (CPU cores)
Yearly read volume in exabytes
CPU cores for analysis and reco
4
3
2
1
0
100000
75000
50000
25000
0
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
Year

Disk IO impact on efficiency

Site	Job eff.	All files	Local files	Remote files
ORNL 2291 jobs (1.223%)	104.6%	19782 files 37.62 MB/s	19512 (98.64%) 39.41 MB/s	270 (1.365%) 6.305 MB/s
Kosice 1216 jobs (0.649%)	70.44%	12178 files 11.45 MB/s	8741 (71.78%) 37.96 MB/s	3437 (28.22%) 4.316 MB/s

Burst read speed

Average over the job duration

Summaries per site

	Number of jobs				PSS			SwapPSS			Average time		CPU	Input Data	
Site	Running	Saving	Done	Error	Min	Avg	Max	Min	Avg	Max	Running ▲	Saving	Efficiency	Size	Rate
ALICE::SaoPaulo::LCG_HTC			1		428.1 MB	428.1 MB	428.1 MB	2.307 GB	2.307 GB	2.307 GB	3:28	32s	57.44%	44.2 GB	3.612 MB/s
ALICE::RAL::LCG			1		2.524 GB	2.524 GB	2.524 GB	2.688 GB	2.688 GB	2.688 GB	2:16	19s	50.33%	59.53 GB	7.419 MB/s
ALICE::Oxford::ARC			1		144.4 MB	144.4 MB	144.4 MB	2.252 GB	2.252 GB	2.252 GB	1:49	33s	52.08%	43.53 GB	6.749 MB/s
ALICE::ORNL::ORNL			6	1	186.7 MB	251.5 MB	535.4 MB	2.187 GB	2.422 GB	2.588 GB	1:08	32s	125.7%	337.1 GB	11.85 MB/s

Summary

IO is/remains a priority

- More storage capacity is welcome
- More spindles in an instance preferred over fragmented capacity
- 10 MB/s/core target

AVX2 will be required

IPv6 on all nodes

Whole node and long TTL queues are better for us

cgroups v2 + lingering enabled