



JAliEn deployment at sites

ALICE Tier-1/Tier-2 Workshop

Budapest

September 26, 2022

v1.0

M. Litmaath

VObos transition recap



- Situation for Run 1 and Run 2
 - Site VOboses run AliEn
 - AliEn submits Job Agents (*pilots*) to the site's computing resources
 - Each Job Agent (JA) runs 1 single-core task at a time
 - If there is no (more) work for a JA, it exits
- Situation for Run 3 and beyond
 - Site VOboses run JAliEn
 - On most VOboses, JAliEn submits 8-core Job Agents
 - For some sites or resources, single-core jobs may need to be used
 - At any time, each 8-core JA can run either of these:
 - One 8-core Run-3 task
 - Up to 8 single-core tasks in parallel, for Run-3 or legacy workflows
 - Each task is run in a Singularity / Apptainer container if possible
 - Sites normally support that these days
 - Standard practice for all 4 LHC experiments and others

Batch system transition



- Most sites need to support 8-core job slots
- Not a problem today: being used by ATLAS and CMS since years
- Configuration depends on the setup of the site
 - ARC CE and HTCondor CE forward requests to their batch system
 - Batch systems may need 8-core queues and/or *fragmentation tuning*
- Current 8-core sites:
 - CERN: all main production “sites” except AURORA → ~1/2 of the capacity
 - T1: CCIN2P3, CNAF, FZK
 - T2/AF: GSI, Hiroshima, NIPNE, Subatech{,_CC IPL}, Vienna, Wigner_KFKI_AF_8core
- 8-core fraction currently is ~40% of the grid capacity
 - To be ramped up in the coming months

JAliEn versions



```
$ ls -l /cvmfs/alice.cern.ch/e16-x86_64/Modules/modulefiles/JAliEn
total 30
-rw-r--r--. 1 cvmfs cvmfs 414 Jun 21 2018 1.0.0-1
[...]
-rw-r--r--. 1 cvmfs cvmfs 453 Sep 22 2021 1.4.1-1
-rw-r--r--. 1 cvmfs cvmfs 453 Oct 6 2021 1.4.2-1
-rw-r--r--. 1 cvmfs cvmfs 446 Nov 4 2021 1.4.3-1
-rw-r--r--. 1 cvmfs cvmfs 446 Nov 12 2021 1.4.4-1
-rw-r--r--. 1 cvmfs cvmfs 446 Nov 18 2021 1.4.4-2
-rw-r--r--. 1 cvmfs cvmfs 446 Dec 13 2021 1.4.5-1
-rw-r--r--. 1 cvmfs cvmfs 446 Jan 20 2022 1.4.6-1
-rw-r--r--. 1 cvmfs cvmfs 446 Feb 22 2022 1.4.7-1
-rw-r--r--. 1 cvmfs cvmfs 446 Mar 3 2022 1.4.7-2
-rw-r--r--. 1 cvmfs cvmfs 446 Mar 21 2022 1.4.8-1
-rw-r--r--. 1 cvmfs cvmfs 446 May 5 19:40 1.4.9-1
-rw-r--r--. 1 cvmfs cvmfs 446 May 18 19:00 1.5.0-1
-rw-r--r--. 1 cvmfs cvmfs 446 May 19 21:31 1.5.1-1
-rw-r--r--. 1 cvmfs cvmfs 446 May 30 17:10 1.5.2-1
-rw-r--r--. 1 cvmfs cvmfs 446 Jun 8 19:10 1.5.3-1
$ ls -l /cvmfs/alice.cern.ch/e17-x86_64/Modules/modulefiles/JAliEn
total 4
-rw-r--r--. 1 cvmfs cvmfs 446 Jul 2 09:40 1.5.4-1
-rw-r--r--. 1 cvmfs cvmfs 446 Jul 10 00:41 1.5.5-1
-rw-r--r--. 1 cvmfs cvmfs 448 Aug 3 16:50 1.5.6-1
-rw-r--r--. 1 cvmfs cvmfs 448 Aug 24 16:20 1.5.7-1
-rw-r--r--. 1 cvmfs cvmfs 447 Sep 21 14:30 1.5.8-1
lrwxrwxrwx. 1 cvmfs cvmfs 7 Sep 21 15:25 new -> 1.5.8-1
lrwxrwxrwx. 1 cvmfs cvmfs 7 Aug 26 18:15 pro -> 1.5.7-1
$
```

- 65 releases already
- On average about 2 releases per month in the last months
- Steady progress with implementing the functionality roadmap
- Plus fixes for issues observed with an ever increasing variety of workloads

JAliEn share



- JAliEn fraction of the resources currently is >99%
 - Remaining 2 sites to be done in the coming weeks
- Reminder: legacy workflows with bad JDLS fail on JAliEn
 - All declared output files *have* to be present for a task to succeed
 - Some failures were still seen earlier this year



MonALISA Repository for ALICE



My jobs

My home dir

LEGO Trains

Hyperloop

Administration Section

Alert XML Feed

JAliEn docs

MonaLisa GUI

The deployment can be tracked on [this page](#)

What is this about?

Status of proxies and AliEn and LCG tests

| Service | AliEn Tests | | | | | | | | | | LCG Tests | | | | | | |
|-------------------|-------------|-----------|--------|-------|-------------|-----------|--------|---------|---------|---------------|--------------|--------|----------------------|--------|-----------|--------|------------|
| | AliEn proxy | | LDAP | | JAliEn cert | | CE | | | Proxy Renewal | Proxy Server | | Proxy of the machine | | Host cert | | |
| | Status | Time left | Status | Cores | Status | Time left | Status | Config | Running | JobAgent | Status | Status | Time left | Status | Time left | Status | Time left |
| 1. Altaria | OK | 1d 23:21 | OK | 1 | OK | 59d 23:48 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:27 | OK | 21:18 | OK | 103d 0:58 |
| 2. Bari_HTC | OK | 1d 23:43 | OK | 1 | OK | 29d 22:57 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:29 | OK | 23:20 | OK | 225d 3:00 |
| 3. Birmingham | OK | 1d 23:20 | OK | 1 | OK | 29d 23:00 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:36 | OK | 22:27 | OK | 135d 18:37 |
| 4. Bratislava_ARC | OK | 1d 23:14 | OK | 1 | OK | 29d 22:48 | OK | new | 1.5.8-1 | 1.5.8-1 | OK | OK | 368d 23:06 | OK | 21:20 | OK | 265d 15:15 |
| 5. Capella | OK | 1d 23:28 | OK | 1 | OK | 60d 0:00 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:35 | OK | 21:25 | OK | 211d 0:36 |
| 6. Catania-VF | OK | 1d 23:15 | OK | 1 | OK | 29d 22:53 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:37 | OK | 21:28 | OK | 263d 3:08 |
| 7. CCIN2P3_HTC | OK | 1d 23:36 | OK | 8 | OK | 29d 22:51 | OK | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 | OK | OK | 201d 12:26 | OK | 21:17 | OK | 249d 10:26 |
| 8. CCIN2P3_HTC_2 | OK | 1d 23:26 | OK | 8 | OK | 29d 22:54 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:37 | OK | 21:27 | OK | 249d 10:37 |
| 9. CERN-AURORA | OK | 1d 23:59 | OK | 1 | OK | 59d 23:51 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:39 | OK | 21:30 | OK | 352d 23:10 |
| 10. CERN-CORONA | OK | 1d 23:30 | OK | 8 | OK | 60d 0:01 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:38 | OK | 21:29 | OK | 211d 0:54 |
| 11. CERN-MIRAGE | OK | 1d 23:16 | OK | 8 | OK | 59d 23:49 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:38 | OK | 21:29 | OK | 189d 0:09 |
| 12. CERN-SIRIUS | OK | 1d 23:55 | OK | 8 | OK | 60d 0:01 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:38 | OK | 21:29 | OK | 211d 14:29 |
| 13. CERN-TRITON | OK | 1d 23:06 | OK | 8 | OK | 59d 23:55 | OK | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 | OK | OK | 201d 12:33 | OK | 21:24 | OK | 352d 23:05 |
| 14. CERN-ZENITH | OK | 1d 23:45 | OK | 8 | OK | 59d 23:51 | OK | new | 1.5.8-1 | 1.5.8-1 | OK | OK | 201d 12:35 | OK | 21:25 | OK | 135d 0:51 |
| 15. Clermont_ARC | OK | 1d 23:42 | OK | 1 | OK | 29d 23:00 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 201d 12:33 | OK | 21:24 | OK | 25d 11:25 |
| 16. CNAF | OK | 1d 23:32 | OK | 8 | OK | 29d 22:52 | OK | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 | OK | OK | 99d 1:09 | OK | 21:29 | OK | 194d 3:09 |
| 17. CNAF-DUE | OK | 1d 23:53 | OK | 8 | OK | 29d 23:02 | OK | pro | 1.5.7-1 | 1.5.7-1 | OK | OK | 99d 1:09 | OK | 21:29 | OK | 41d 3:09 |

JAliEn versions at sites



| Service | AliEn proxy | | LDAP | | JAliEn cert | | CE | | | |
|----------------------|-------------|-----------|--------|-------|-------------|-----------|--------|---------|---------|------------|
| | Status | Time left | Status | Cores | Status | Time left | Status | Config | Running | JobAgent ▲ |
| 75. Yerevan | | 1d 23:06 | | 1 | | 60d 0:25 | | pro | local | custom |
| 7. CCIN2P3_HTC | | 1d 23:06 | | 8 | | 29d 23:21 | | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 |
| 13. CERN-TRITON | | 1d 23:36 | | 8 | | 60d 0:25 | | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 |
| 16. CNAF | | 1d 23:02 | | 8 | | 29d 23:22 | | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 |
| 22. FZK | | 1d 23:19 | | 8 | | 29d 23:21 | | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 |
| 54. RAL | | 1d 23:30 | | 1 | | 29d 23:31 | | 1.5.8-1 | 1.5.8-1 | 1.5.8-1 |
| 4. Bratislava_ARC | | 1d 23:45 | | 1 | | 29d 23:18 | | new | 1.5.8-1 | 1.5.8-1 |
| 14. CERN-ZENITH | | 1d 23:15 | | 8 | | 60d 0:21 | | new | 1.5.8-1 | 1.5.8-1 |
| 25. GRIF_IPNO_IJCLAB | | 1d 23:40 | | 1 | | 17d 2:24 | | new | 1.5.8-1 | 1.5.8-1 |
| 47. Oxford | | 1d 23:28 | | 1 | | 29d 23:25 | | new | 1.5.8-1 | 1.5.8-1 |
| 63. Subatech | | 1d 23:25 | | 8 | | 29d 23:27 | | new | 1.5.8-1 | 1.5.8-1 |
| 71. Vienna | | 1d 23:00 | | 8 | | 29d 23:36 | | new | 1.5.8-1 | 1.5.8-1 |
| 1. Altaría | | 1d 23:51 | | 1 | | 60d 0:18 | | pro | 1.5.7-1 | 1.5.7-1 |
| 2. Bari_HTC | | 1d 23:13 | | 1 | | 29d 23:27 | | pro | 1.5.7-1 | 1.5.7-1 |
| 3. Birmingham | | 1d 23:50 | | 1 | | 29d 23:30 | | pro | 1.5.7-1 | 1.5.7-1 |
| 5. Capella | | 1d 23:59 | | 1 | | 60d 0:30 | | pro | 1.5.7-1 | 1.5.7-1 |
| 6. Catania-VF | | 1d 23:45 | | 1 | | 29d 23:23 | | pro | 1.5.7-1 | 1.5.7-1 |
| 8. CCIN2P3_HTC_2 | | 1d 23:56 | | 8 | | 29d 23:24 | | pro | 1.5.7-1 | 1.5.7-1 |
| 9. CERN-AURORA | | 1d 23:29 | | 1 | | 60d 0:22 | | pro | 1.5.7-1 | 1.5.7-1 |
| 10. CERN-CORONA | | 1d 23:00 | | 8 | | 60d 0:31 | | pro | 1.5.7-1 | 1.5.7-1 |
| 11. CERN-MIRAGE | | 1d 23:46 | | 8 | | 60d 0:19 | | pro | 1.5.7-1 | 1.5.7-1 |
| 12. CERN-SIRIUS | | 1d 23:25 | | 8 | | 60d 0:31 | | pro | 1.5.7-1 | 1.5.7-1 |
| 15. Clermont_ARC | | 1d 23:11 | | 1 | | 29d 23:30 | | pro | 1.5.7-1 | 1.5.7-1 |
| 17. CNAF-DUE | | 1d 23:23 | | 8 | | 29d 23:32 | | pro | 1.5.7-1 | 1.5.7-1 |

Development site

Fixed: final testing

New: early testing

Default production

VObbox perspective recap



- AliEn services
 - **CE**: submits JAs up to a limit when there is compatible work
 - **ClusterMonitor**: receives heartbeat messages from the JAs
 - **CMreport**: forwards those messages to the central services
 - **MonALISA**: monitors jobs, services, SE and network
 - Needed some trickery to start automatically at boot time (which could fail)
 - **Commands**: *alien*, *aliend*
- JAliEn services
 - **CE**: now needs an extra credential to obtain job tokens for JAs
 - \$ openssl x509 -noout -subject -in .globus/hostcert.pem
 - subject= /C=ch/O=AliEn2/OU=ALICE/CN=voboxalice10.cern.ch
 - **MonALISA**
 - JA heartbeat messages go directly to the central services
 - Crontab entries restart the services as needed, with the right environment
 - **Command**: */cvmfs/alice.cern.ch/scripts/vobox/jalien-vobox.sh*
 - Thanks to Kalana Wijethunga for the design and the first versions!

VBox configuration



```
$ cat .alien/config/CE.env
# CE environment bootstrap

export X509_USER_PROXY=$(
    ls -t /var/lib/vobox/alice/proxy_repository/*lcgadmin | sed q
)

$ cat .alien/config/ml.env
# MonaLisa environment bootstrap

export X509_USER_PROXY=$(
    ls -t /var/lib/vobox/alice/proxy_repository/*lcgadmin | sed q
)
export MYPROXY_SERVER=myproxy.cern.ch

$ cat .alien/config/version.properties
#jalien=x.y.z-r
#MonaLisa=yyyymmdd-r

$ cat .j/config/version.properties
#jobagent.version=x.y.z-r
#custom.jobagent.jar.dir=/cvmfs/alice.cern.ch/java/jar/jalien
```

- Define X509 proxy for CE and ML as needed
- Specify CE or ML version as needed
 - Defaults normally OK
 - New versions tested at selected sites
- Specify JAliEn version to be used by the JA
 - By default the version used by the CE, normally OK

From X509 to WLCG tokens



- WLCG Authorization WG is coordinating a transition from X509 + VOMS to WLCG tokens during Run 3 ([link](#))
 - Inspired by the use of tokens in ALICE since many years!
 - And by common practice in industry and academia
 - To support federated identities and token standards
- VO management will go from VOMS-Admin to (INDIGO-) IAM
 - Identity and Access Management
 - Already being used in production by CMS and ATLAS (see [July GDB talk](#))
 - It can provide finely scoped tokens to trusted entities
 - It also has a VOMS endpoint for backward compatibility
- The first deadline that concerns us here is Feb 2023: EOL for GSI (X509 proxy) authentication in the HTCondor CE
 - The JAliEn CE is ready for using WLCG tokens – waiting for the ALICE IAM
 - A token renewal service will provide VOboxes with those securely
 - Already in production for JAliEn CE credential renewals

[WLCG Token Transition Timeline v1.0](#) published on August 22 on Zenodo