

JAliEn Dev & Ops

costin.grigoras@cern.ch

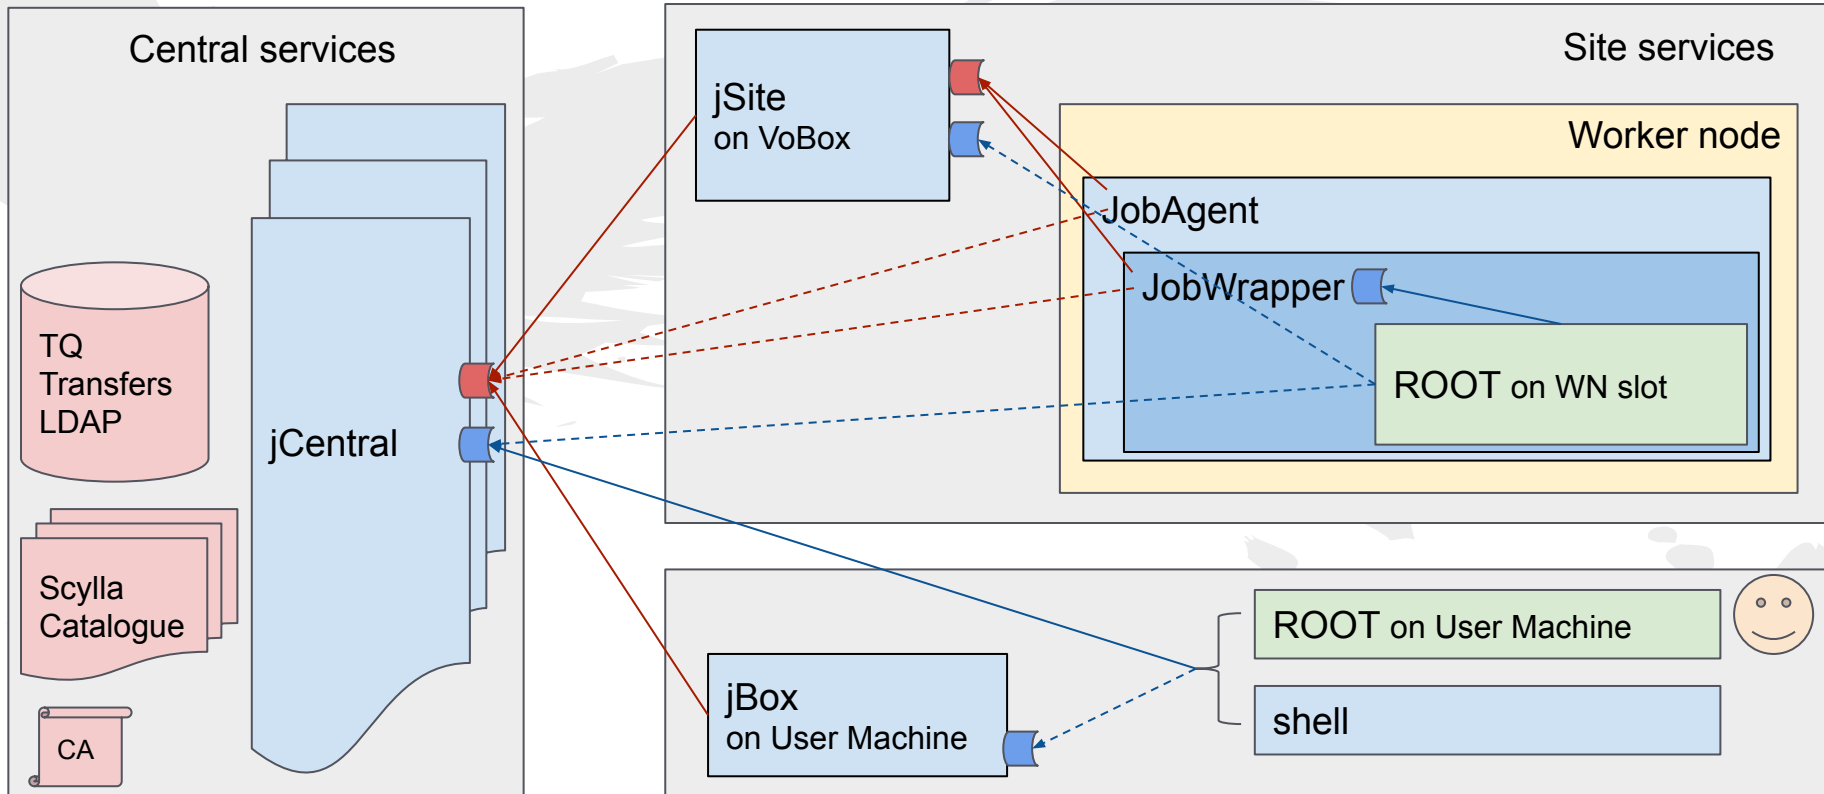
JAliEn

—→ Default uplink

🔴 SSL(Compressed(Java serialized object stream)) (TCP/8098)

- - - - - Optional uplink

🔵 WebSocketS, JSON serialization of requests/replies (TCP/8097)



Central services

Deployed on 7 of the central machines

```
$ host alice-jcentral.cern.ch  
alice-jcentral.cern.ch has address 137.138.47.216  
alice-jcentral.cern.ch has address 137.138.47.221  
alice-jcentral.cern.ch has address 137.138.99.136  
alice-jcentral.cern.ch has address 137.138.99.139  
alice-jcentral.cern.ch has address 137.138.99.140  
alice-jcentral.cern.ch has address 137.138.99.141  
alice-jcentral.cern.ch has address 137.138.99.145  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b49f::100:2  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b49f::100:5  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b49f::100:6  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b49f::100:7  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b49f::100:27  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b50e::100:12  
alice-jcentral.cern.ch has IPv6 address 2001:1458:201:b50e::100:17
```

Starting from CVMFS

```
[grigoras@lxplus ~]$ alienv enter JAliEn
```

```
[JAliEn] ~ > jalien
```

```
[Grid certificate password: ]
```

```
jsh:[alice] [01] /alice/cern.ch/user/g/grigoras/ >
```

← First execution loads the identity

```
[JAliEn] ~ > ps a | grep JBox
```

```
grigoras ... java -client -Xmx64m -Xms64m -Djava.io.tmpdir=/tmp/grigoras -Duserid=2010 alien.JBox
```

```
[JAliEn] ~ > jalien
```

```
jsh:[alice] [01] /alice/cern.ch/user/g/grigoras/ >
```

← Subsequent calls find the running agent

```
[JAliEn] ~ > ls -lh $CLASSPATH
```

```
-rw-r--r--. 1 cvmfs cvmfs 14M .../cvmfs/alice.cern.ch/.../JAliEn/1.0.9-1/lib/ alien-users.jar
```

One JAR file packing all necessary libraries

```
[JAliEn] ~ > java -version
```

```
openjdk version "10.0.1" 2018-04-17
```

```
[JAliEn] ~ > xrdcp -version
```

```
v4.8.5
```

Recent versions of the two external dependencies

Improved shell cmd: *mirror*

```
jsh:[alice] [01] /alice/cern.ch/user/g/grigoras/ > whereis JDL
the file JDL is in
  SE => ALICE::SARA::DCACHE
  SE => ALICE::CLERMONT::SE
  SE => ALICE::CNAF::SE
  SE => ALICE::ZA_CHPC::EOS

jsh:[alice] [02] /alice/cern.ch/user/g/grigoras/ > mirror -S disk:4 JDL
ALICE::CERN::EOS: queued transfer ID 221777638
ALICE::FZK::SE: queued transfer ID 221777639
ALICE::RRC_KI_T1::EOS: queued transfer ID 221777640
ALICE::PRAGUE::SE: queued transfer ID 221777641

jsh:[alice] [05] /alice/cern.ch/user/g/grigoras/ > whereis JDL
...
  SE => ALICE::ZA_CHPC::EOS
  SE => ALICE::CERN::EOS
  SE => ALICE::RRC_KI_T1::EOS

jsh:[alice] [08] /alice/cern.ch/user/g/grigoras/ > whereis JDL
the file JDL is in
...
  SE => ALICE::RRC_KI_T1::EOS
  SE => ALICE::PRAGUE::SE
  SE => ALICE::FZK::SE
```

Quickly create more copies of
critical files

New *cp* feature

A much faster way to download the internet (multiple threads, skipping existing files to resume/retry)

```
jsh:[alice] [11] /alice/sim/2018/LHC18k9a/276556/AOD/ > cp -T 32 */aod_archive.zip file:/tmp/  
Downloaded file to /tmp/001/aod_archive.zip  
Downloaded file to /tmp/002/aod_archive.zip  
Downloaded file to /tmp/004/aod_archive.zip  
Downloaded file to /tmp/003/aod_archive.zip
```

```
jsh:[alice] [12] /alice/sim/2018/LHC18k9a/276556/AOD/ > cp -T 32 */aod_archive.zip file:/tmp/  
Error: Local copy target /tmp/001/aod_archive.zip exists, skipping it  
Error: Local copy target /tmp/003/aod_archive.zip exists, skipping it  
Error: Local copy target /tmp/002/aod_archive.zip exists, skipping it  
Error: Local copy target /tmp/004/aod_archive.zip exists, skipping it
```

New command: *xrdstat*

```
jsh:...> xrdstat split.jdl
Checking the replicas of /alice/cern.ch/user/g/grigoras/split.jdl
ALICE::CERN::EOS      root://eosalice.cern.ch:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK
ALICE::ORNL::EOS      root://alice-eos-01.ornl.gov:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK
ALICE::CCIN2P3::SE    root://ccxrdralice.in2p3.fr:1096//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      ERR
ALICE::Legnaro::SE    root://t2-xrdrd.lnl.infn.it:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK

jsh:...> xrdstat -d -v split.jdl
ALICE::CERN::EOS      root://eosalice.cern.ch:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK
Downloaded file matches the catalogue details, retrieving took 0s (1.571 KB/s)
ALICE::CCIN2P3::SE    root://ccxrdralice.in2p3.fr:1096//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      ERR
Command has timed out and was killed after 15s, retry #1, output was , for command : [xrdfs, ccxrdralice.in2p3.fr:1096, stat, /02/05607/02e378b0-8d45-11e5-a935-2c44fd849358]
ALICE::Legnaro::SE    root://t2-xrdrd.lnl.infn.it:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK
Downloaded file matches the catalogue details, retrieving took 0s (1.867 KB/s)
ALICE::ORNL::EOS      root://alice-eos-01.ornl.gov:1094//02/05607/02e378b0-8d45-11e5-a935-2c44fd849358      OK
Downloaded file matches the catalogue details, retrieving took 1s (214.2 B/s)

jsh:...> xrdstat -p 1475408332
Checking the replicas of /alice/data/2018/LHC18q/000296244/pass1/18000296244039.5005/root_archive.zip
ALICE::Bratislava::SE  root://lcgstorage04.dnp.fmph.uniba.sk:1094//02/59097/d917bd90-286c-11e9-87ae-c3b9ca074496      OK
ALICE::CERN::EOS      root://eosalice.cern.ch:1094//02/59097/d917bd90-286c-11e9-87ae-c3b9ca074496      OK
Checking the replicas of /alice/data/2018/LHC18q/000296244/pass1/18000296244039.501/root_archive.zip
ALICE::Bratislava::SE  root://lcgstorage04.dnp.fmph.uniba.sk:1094//05/10013/9bbddc14-2870-11e9-bf77-a76507f7d16c      OK
ALICE::CERN::EOS      root://eosalice.cern.ch:1094//05/10013/9bbddc14-2870-11e9-bf77-a76507f7d16c      OK
Checking the replicas of /alice/data/2018/LHC18q/000296244/pass1/18000296244039.5013/root_archive.zip
ALICE::Bratislava::SE  root://lcgstorage04.dnp.fmph.uniba.sk:1094//08/53690/43f3a0f8-286c-11e9-a6c1-c7a53c9b7b23      ERR
ALICE::CERN::EOS      root://eosalice.cern.ch:1094//08/53690/43f3a0f8-286c-11e9-a6c1-c7a53c9b7b23      OK
```

New command: *w*

```
jsh:... > w
```

Account name	Active jobs	Waiting jobs
1. alidaq	8530	39
2. aliproduct	82870	54027
3. alitrain	17774	7369
...		
56. *****	41	15
TOTAL	119010	63892

New command: *listSEs*

```
jsh:[... > listSEs -qos tape
```

SE name	ID	Total	Capacity			Demote			Endpoint URL
			Used	Free		Read	Write	QoS	
ALICE::CCIN2P3::TAPE	306	327.4 TB	2.516 PB	0 B		0.0000	1.3189	tape	root://ccxrdralice.in2p3.fr:1094
ALICE::CERN::T0ALICE	273	3.352 PB	34.61 PB	0 B		0.0330	0.4883	tape	root://castoralice.cern.ch:1094
ALICE::CNAF::TAPE	295	436.1 TB	9.432 PB	0 B		0.0000	0.2673	tape	root://alice-xrootd-tsm.cr.cnaf.infn.it:
ALICE::FZK::TAPE	286	1.175 PB	7.664 PB	0 B		0.1453	1.3046	tape	root://alice-tape-se.gridka.de:1094

```
...
```

New command: *listSEDistance*

```
jsh:... > listSEDistance -qos disk -site Hiroshima
```

```
ALICE::HIROSHIMA::EOS      (read: 0.000, write: -0.164, distance: -0.164)
ALICE::KISTI_GSDC::SE2    (read: 0.000, write: -0.288, distance: -0.068)
ALICE::FZK::SE            (read: 0.007, write: -0.506, distance: -0.063)
ALICE::CERN::EOS          (read: 0.000, write: -0.601, distance: -0.010)
ALICE::RRC_KI_T1::EOS     (read: 0.000, write: -0.475, distance: 0.095)
```

```
jsh:... > listSEDistance -read split.jdl
```

```
ALICE::CERN::EOS          (read: 0.000, write: -0.601, distance: 0.004)
ALICE::CCIN2P3::SE       (read: 0.001, write: 1.024, distance: 0.007)
ALICE::LEGNARO::SE       (read: 0.000, write: -0.122, distance: 0.031)
ALICE::ORNL::EOS         (read: 1.000, write: 2.008, distance: 1.214)
```


Client code

ROOT plugin for AliEn refactored

- Legacy code extracted from ROOT
- Build recipes for both legacy and new plugins
- Parallel builds of both versions

Similar to the ROOT5 -> ROOT6 transition

- Switched to Xrootd 4.* libraries
- In testing phase now

More clients

Command line Python scripts

Adrian will give you the details on them

Ådne is expanding the framework to implement a full shell

Using the WebSocket(S) interface to connect upstream and Xrootd python bindings for efficient IO

WebSocket interface

Server side implementation ready

Running the binary-to-JSON translation as close to the client as possible

JobWrapper, VoBox services, CS - all expose it

Intermediate services communicate upstream on the Java compressed and encrypted object stream

Firewall openings

WNs to Central Services and VoBox:

TCP/8097 (WebSocketS)

TCP/8098 (Java Object Stream)

CS are currently on these subnets:

IPv4: 137.138.99.128/26, 137.138.47.192/26,
188.184.2.0/26

IPv6: 2001:1458:201:22::/64, 2001:1458:201:b49f::/64,
2001:1458:201:b50e::/64

VoBox port summary

Full details [here](#)

- 8098/TCP incoming from **site WN** - JAliEn/Java Serialized Object stream
 - 8097/TCP incoming from **site WN** - JAliEn/WebSocketS
 - 8084/TCP incoming from **CERN** and the **site WN** - ClusterMonitor
 - 1093/TCP incoming from World - MonALISA FDT server, SE tests
 - 8884/UDP incoming from the **site WN** and **site SE nodes** - Monitoring info
 - 9930/UDP incoming from the **site SE nodes** - Xrootd metrics
 - + **ICMP** incoming and outgoing - network topology for file placement and access
- } Two new ports

IPv6 status

Please make sure the previous firewall openings are applied on both protocols,

as well as having default routes, forward and reverse DNS, etc...

All new services are dual stacked and prefer IPv6 when available

New IPv6-only storage tests

Storage IPv6 status

By count:
~50% of instances

By volume:
61% disk total space
28% tape content

IPv6 deployment on SEs

