



EUROPEAN MIDDLEWARE INITIATIVE

New Data Staging in the ARC CE Tutorial

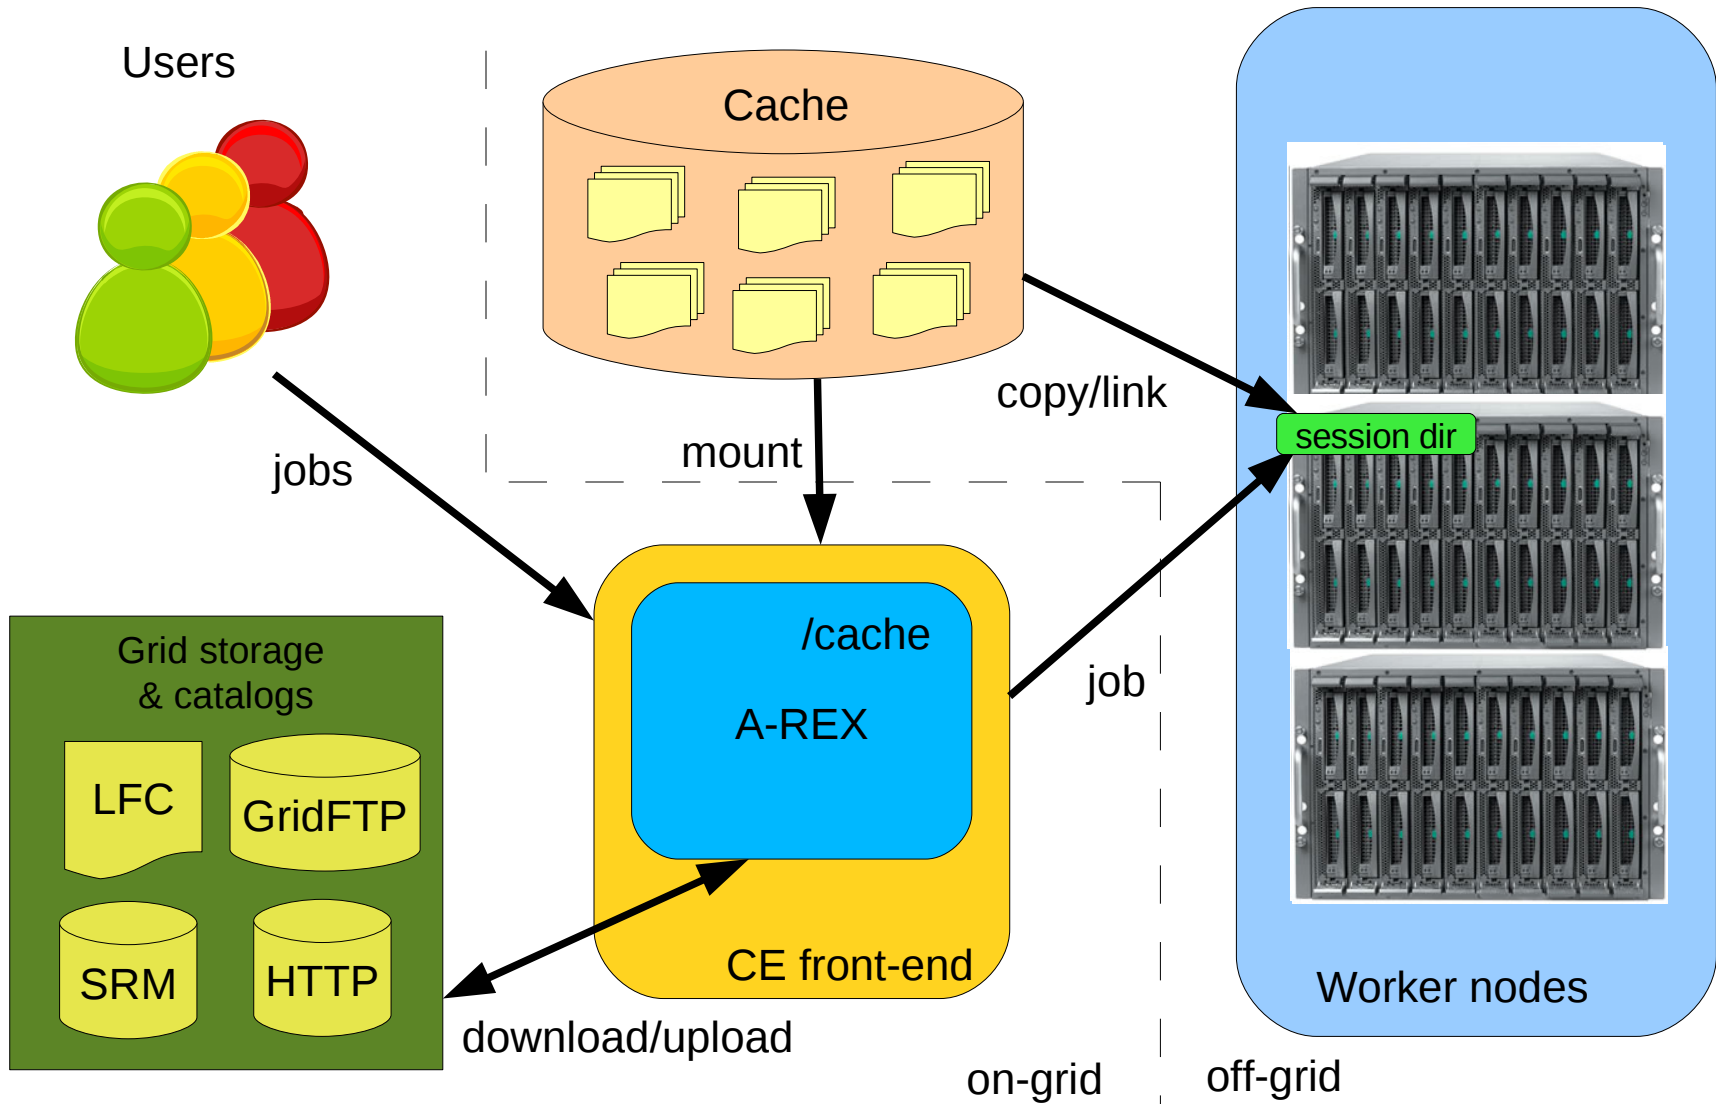
David Cameron
University of Oslo

Introduction

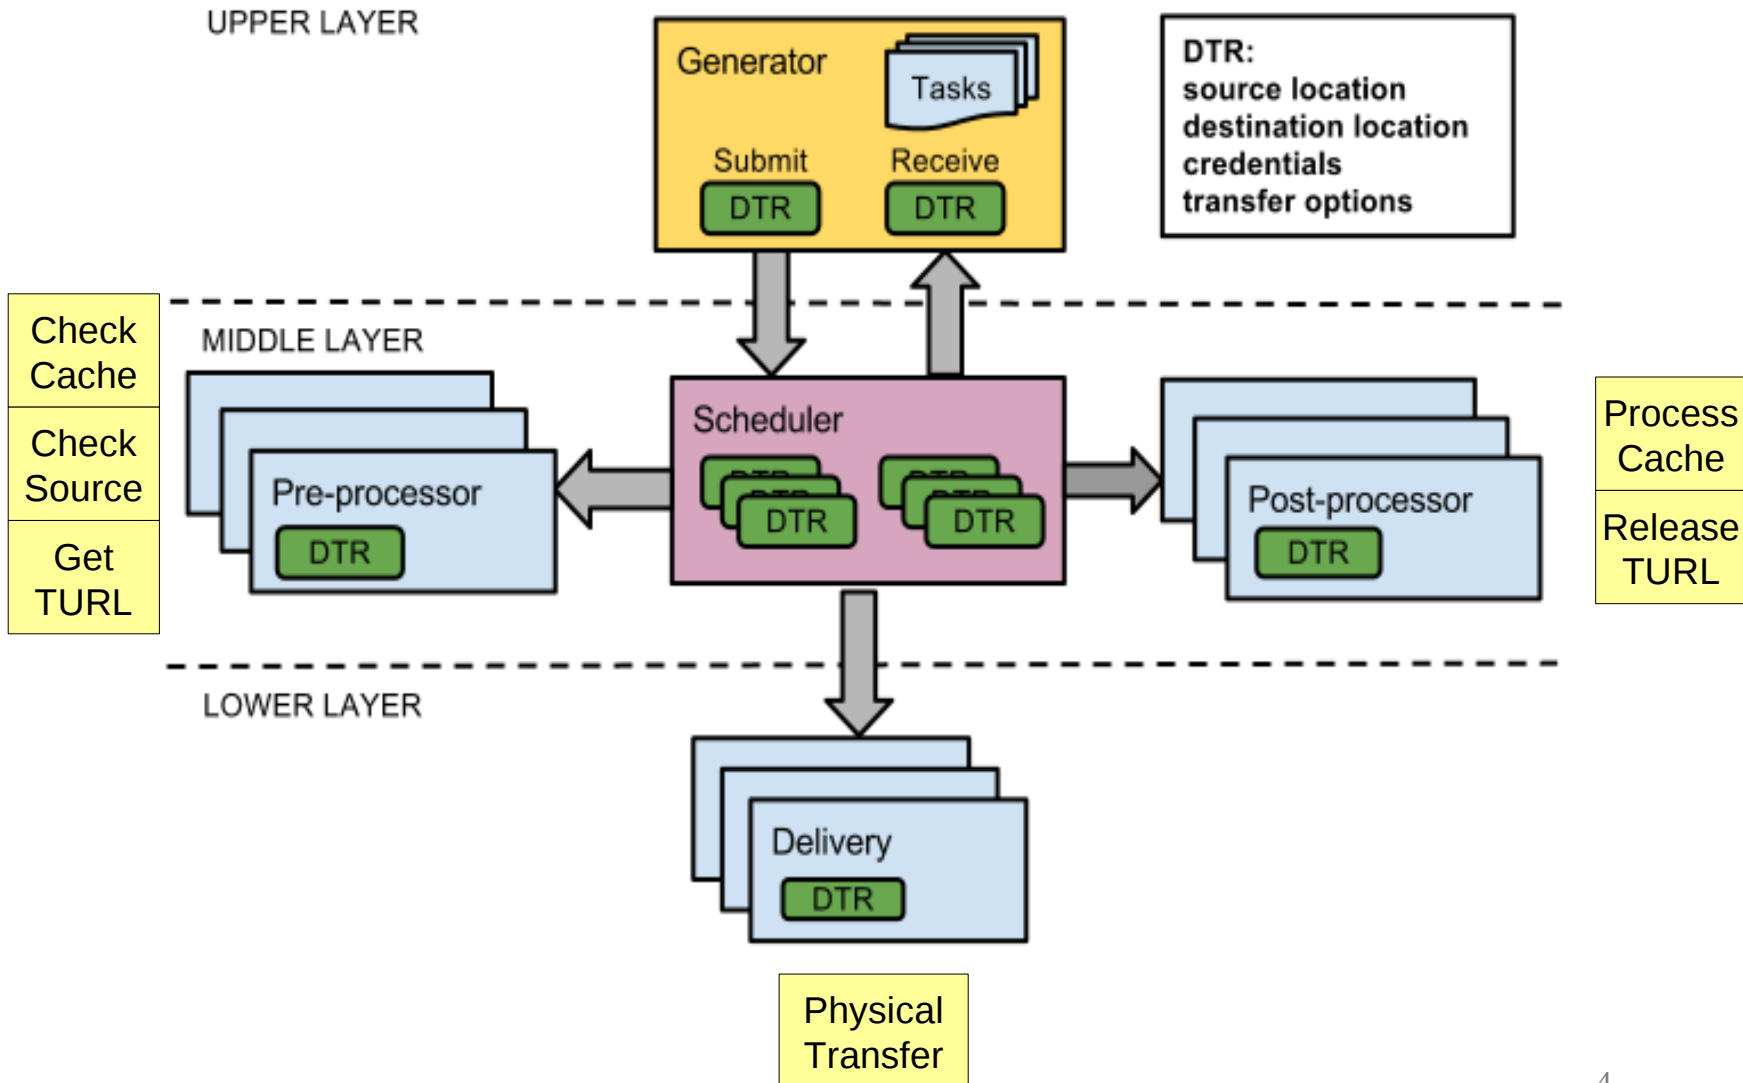


- This tutorial is aimed at sites running (or interested in running) ARC
 - For more internal technical details see talk on Tuesday
 - <https://www.egi.eu/indico/contributionDisplay.py?sessionId=52&contribId=151&confId=679>
- Covers configuration, deployment and monitoring

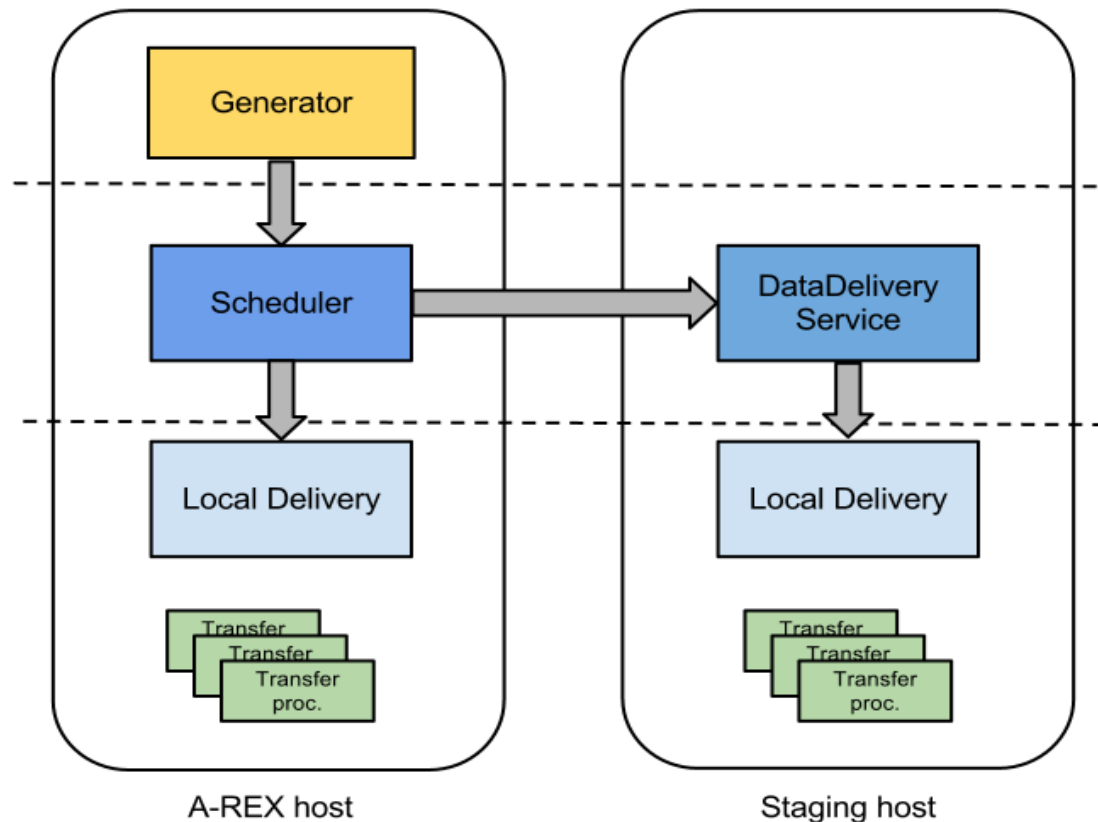
ARC Data Management Architecture



Design: Data Transfer Request



Multi-host Staging



- Delivery layer can be expanded to multiple hosts
- Data delivery service allows Scheduler to launch transfers on remote hosts

What does it mean?

- Intelligent queuing system for data transfer
- Efficient use of network bandwidth
- Cache files can be processed instantly
- Priority and fair-share system
- Easily expandable to multiple hosts

Setup and Configuration: A-REX

- If the ARC CE (A-REX) is installed, nothing else is needed
 - A configuration option in arc.conf turns on the new system
 - newdatastaging=yes
 - Many parameters can be configured:
 - http://wiki.nordugrid.org/index.php/Data_Staging

```
[data-staging]
maxdelivery="10"
maxprocessor="20"
maxemergency="2"
maxprepared="50"
```

```
sharetype="voms:role"
definedshare="myvo:production 80"
```

```
deliveryservice="https://spare.host:60003/datadeliveryservice"
```

Setup and Configuration: Remote Hosts

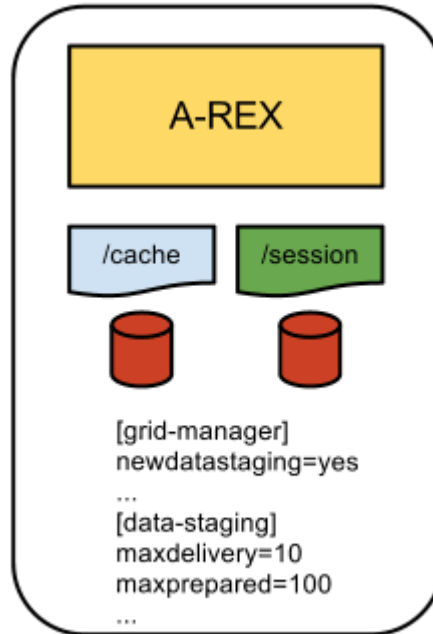
- On each host install **nordugrid-arc-datadelivery-service** (and get a host certificate)
- IMPORTANT! Paths must be the same on all hosts
- Configure allowed directories and other options
 - http://wiki.nordugrid.org/index.php/Data_Staging/Multi-host

```
[common]
loglevel = INFO
interface = myhost.org
port = 60002
```

```
[datadelivery-service]
allowed_ip = 1.2.3.4
allowed_dir = /var/arc/session
allowed_dir = /var/arc/cache
```

- Start service
 - `/etc/init.d/arc-datadelivery-service start`

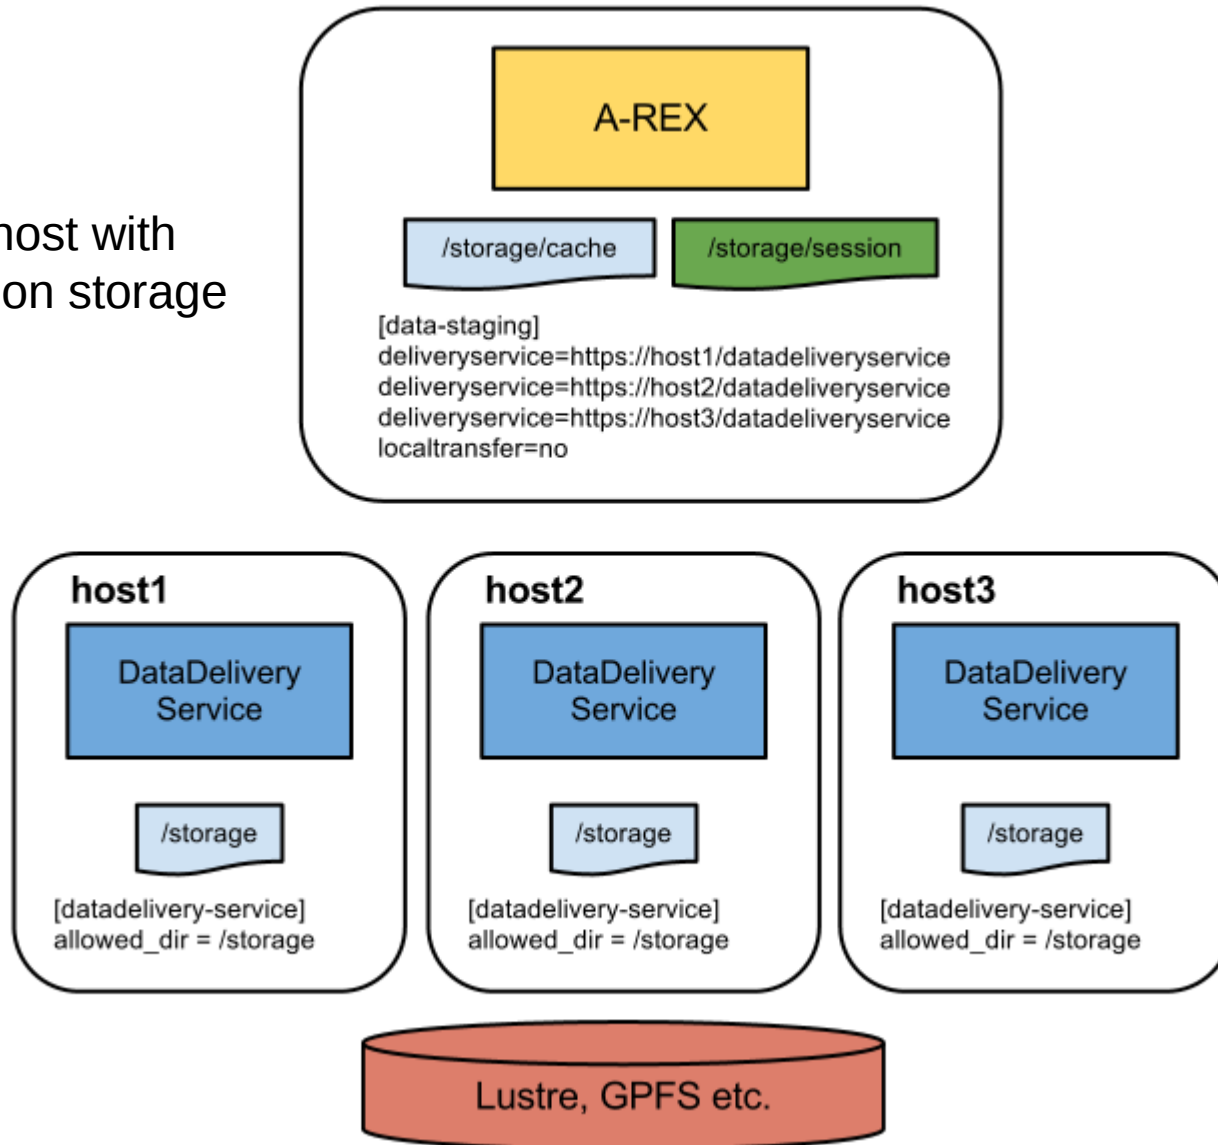
Deployment: Small site



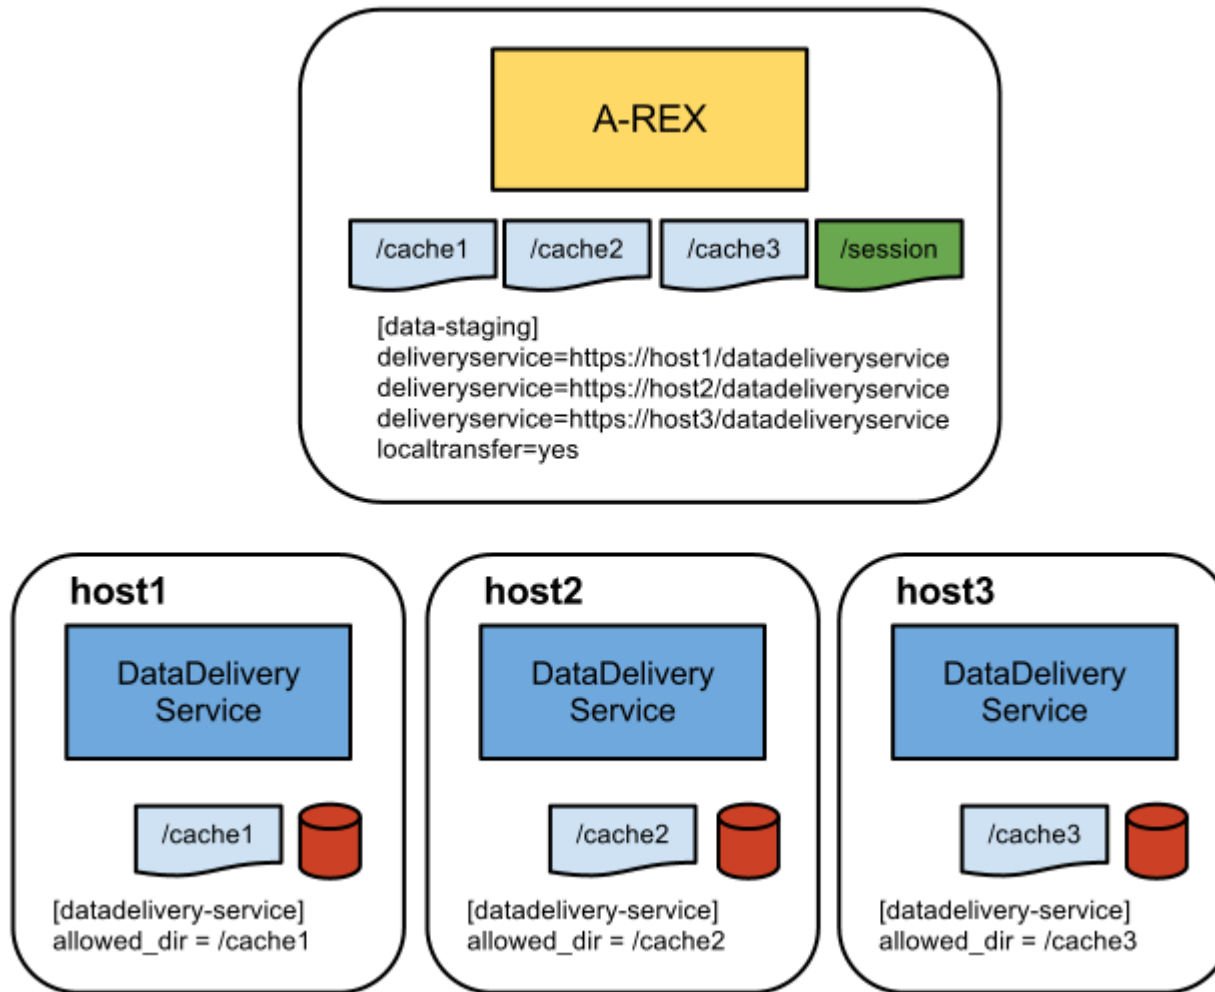
Single CE node with local or NFS-mounted cache and session

Deployment: Large Site with Shared Storage

Multi-host with
common storage



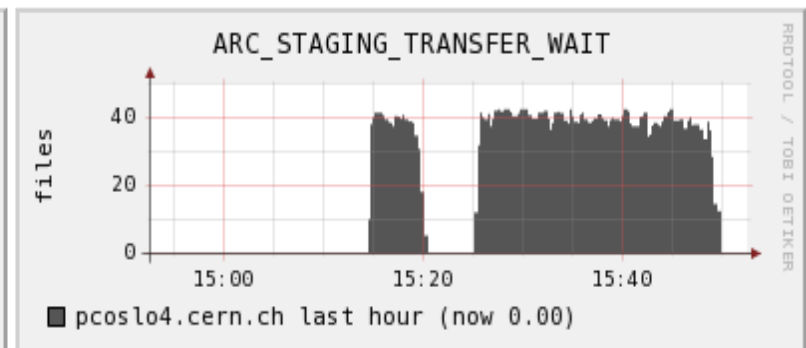
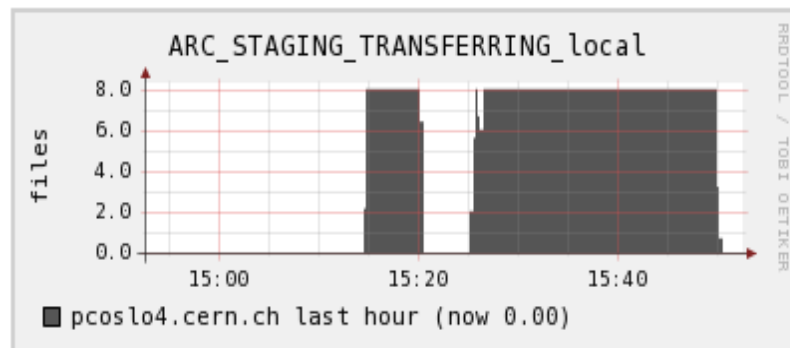
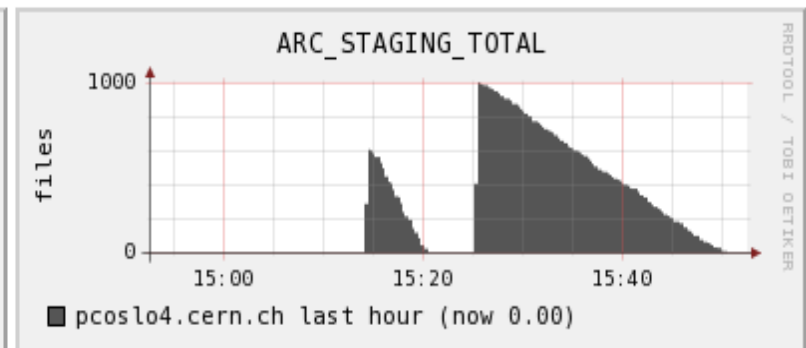
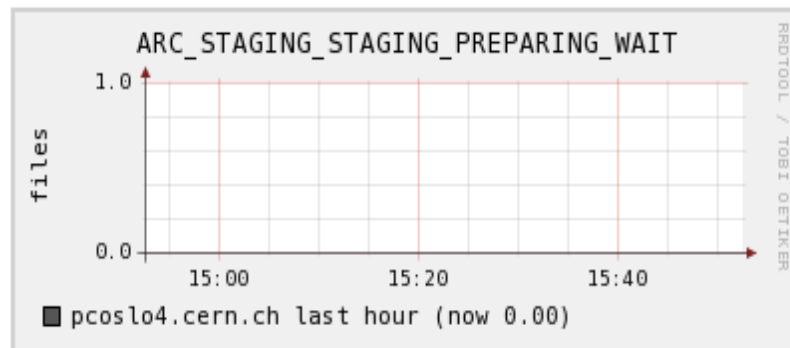
Deployment: Large Site with Local Caches



Multi-host with local caches

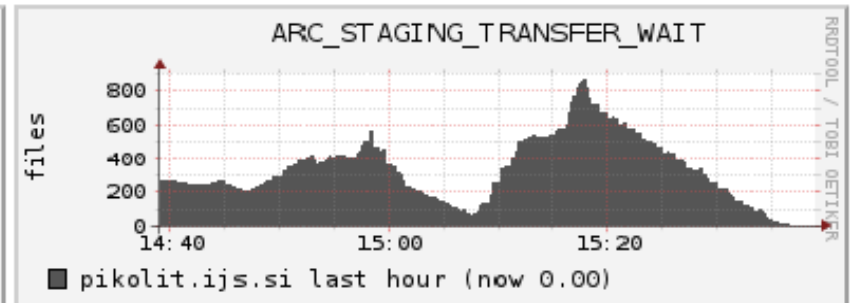
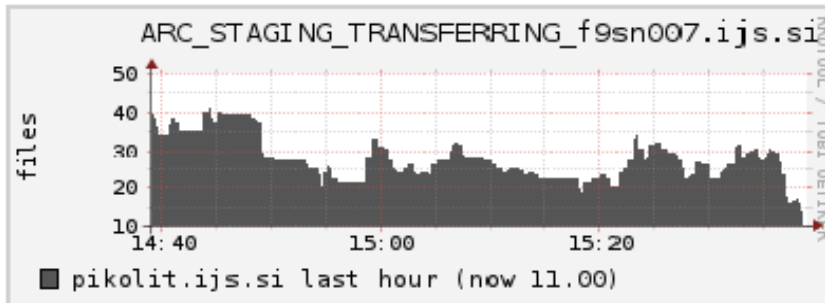
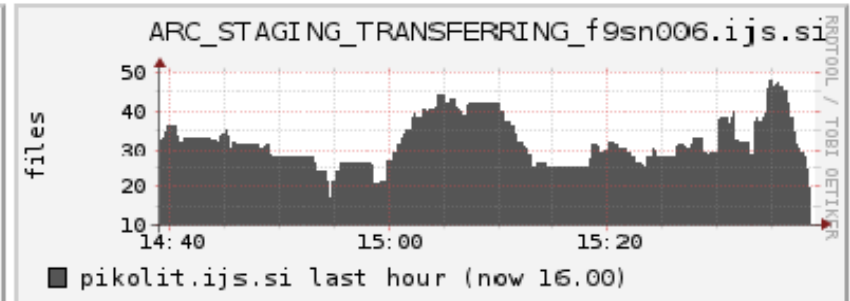

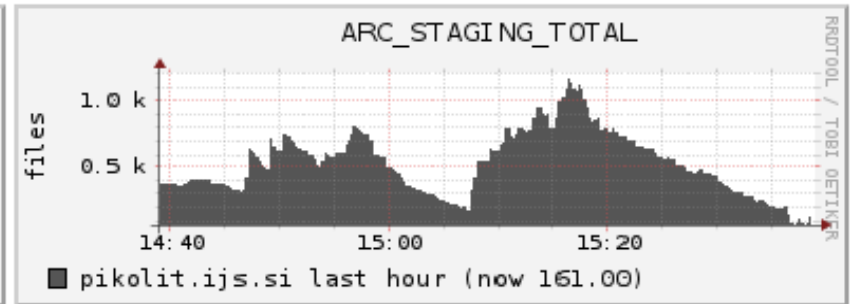

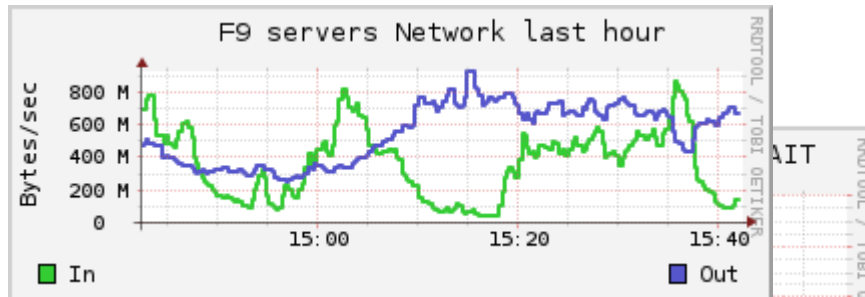
Monitoring

- Graphical monitoring via Gangliarc
 - Adds ARC-related ganglia metrics via gmetric
 - <http://wiki.nordugrid.org/index.php/Gangliarc>



Real Life Usage – ATLAS production/analysis

Thanks to Andrej Filipcic
(Jozef Stefan Institute)



Use Case

- Submit a set of 50 jobs requiring input data
 - 10 ATLAS files (~50MB each)
 - Same transfer share
 - Dummy executable
- Half the jobs have priority 50, the other half 100 (randomly selected)

Snapshot

> arcstat -a

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/Vm4NDmEFLVfncsfcko13THIpABFKDmABFKDmu6JKDmABFKDmrz9Xrn
Name: random-p100
State: Finished (FINISHED)
Exit Code: 0

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/lflKDMFFLVfncsfcko13THIpABFKDmABFKDm36JKDmABFKDm7LEc6n
Name: random-p50
State: Preparing (PREPARING)

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/aZULDmFFLVfncsfcko13THIpABFKDmABFKDmF7JKDmABFKDm7LEc6n
Name: random-p100
State: Finished (FINISHED)
Exit Code: 0

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/r2GMDmFFLVfncsfcko13THIpABFKDmABFKDmO7JKDmABFKDm7LEc6n
Name: random-p50
State: Preparing (PREPARING)

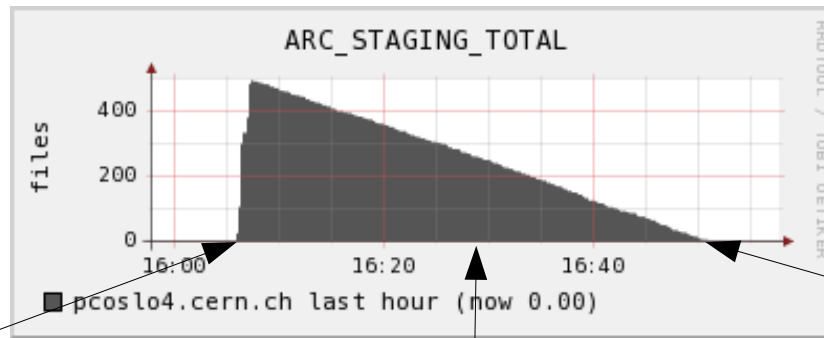
Job: gsiftp://pcoslo4.cern.ch:2811/jobs/S9DNDmFFLVfncsfcko13THIpABFKDmABFKDmY7JKDmABFKDm7LEc6n
Name: random-p100
State: Finished (FINISHED)
Exit Code: 0

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/lvtNDmFFLVfncsfcko13THIpABFKDmABFKDmi7JKDmABFKDm7LEc6n
Name: random-p50
State: Preparing (PREPARING)

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/19PKDmGFLVfncsfcko13THIpABFKDmABFKDmr7JKDmABFKDmVvap9n
Name: random-p50
State: Preparing (PREPARING)

Job: gsiftp://pcoslo4.cern.ch:2811/jobs/oE6KDmGFLVfncsfcko13THIpABFKDmABFKDm37JKDmABFKDmVvap9n
Name: random-p100
State: Finished (FINISHED)
Exit Code: 0

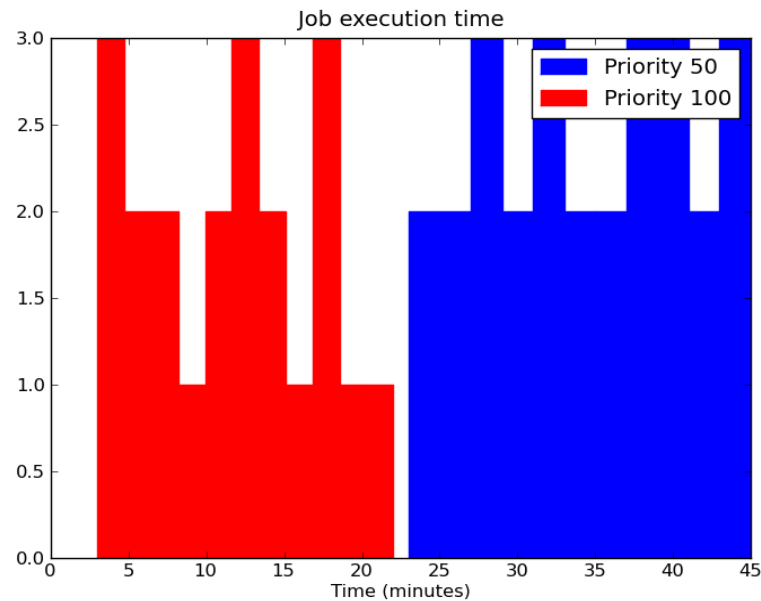
Results



Job submission

Last high priority job finished

Last job finished



Links

- Main wiki page
 - http://wiki.nordugrid.org/index.php/Data_Staging
- Multi-host data staging
 - http://wiki.nordugrid.org/index.php/Data_Staging/Multi-host
- Gangliarc
 - <http://wiki.nordugrid.org/index.php/Gangliarc>
- Talk from Tuesday with technical detail
 - <https://www.egi.eu/indico/contributionDisplay.py?contribId=151&con>



Thank you!

EMI is partially funded by the European Commission under Grant Agreement RI-261611